arXiv:1601.07277v1 [math.OC] 27 Jan 2016

A General System for Heuristic Solution of Convex Problems over Nonconvex Sets

Steven Diamond Reza Takapoui Stephen Boyd

January 28, 2016

Abstract

We describe general heuristics to approximately solve a wide variety of problems with convex objective and decision variables from a nonconvex set. The heuristics, which employ convex relaxations, convex restrictions, local neighbor search methods, and the alternating direction method of multipliers (ADMM), require the solution of a modest number of convex problems, and are meant to apply to general problems, without much tuning. We describe an implementation of these methods in a package called NCVX, as an extension of CVXPY, a Python package for formulating and solving convex optimization problems. We study several examples of well known nonconvex problems, and show that our general purpose heuristics are effective in finding approximate solutions to a wide variety of problems.

1 Introduction

1.1 The problem

We consider the optimization problem

$$\begin{array}{ll} \text{minimize} & f_0(x,z) \\ \text{subject to} & f_i(x,z) \leq 0, \quad i = 1,\ldots,m \\ & Ax + Bz = c \\ & z \in \mathcal{C}, \end{array} \tag{1}$$

where $x \in \mathbf{R}^n$ and $z \in \mathbf{R}^q$ are the decision variables, $A \in \mathbf{R}^{p\times n}$, $B \in \mathbf{R}^{p\times q}$, $c \in \mathbf{R}^p$ are problem data, and $\mathcal{C} \subseteq \mathbf{R}^q$ is compact. We assume that the objective and inequality constraint functions $f_0,\ldots,f_m : \mathbf{R}^n \times \mathbf{R}^q \to \mathbf{R}$ are jointly convex in x and z. When the set $\mathcal{C}$ is convex, (1) is a convex optimization problem, but we are interested here in the case where $\mathcal{C}$ is not convex. Roughly speaking, the problem (1) is a convex optimization problem, with some additional nonconvex constraints, $z \in \mathcal{C}$. We can think of x as the collection of decision variables that appear only in convex constraints, and z as the decision variables that are directly constrained to lie in the (generally) nonconvex set $\mathcal{C}$. The set $\mathcal{C}$

is often a Cartesian product, $\mathcal{C} = \mathcal{C}_1 \times \cdots \times \mathcal{C}_k$, where $\mathcal{C}_i \subset \mathbf{R}^{q_i}$ are sets that are simple to describe, *e.g.*, $\mathcal{C}_i = \{0, 1\}$. We denote the optimal value of the problem (1) as $p^\star$, with the usual conventions that $p^\star = +\infty$ if the problem is infeasible, and $p^\star = -\infty$ if the problem is unbounded below.

1.2 Special cases

Mixed-integer convex optimization. When $\mathcal{C} = \{0, 1\}^q$, the problem (1) is a general mixed integer convex program, *i.e.*, a convex optimization problem in which some variables are constrained to be Boolean. (Mixed Boolean convex program would be a more accurate name for such a problem, but 'mixed integer' is commonly used.) It follows that the problem (1) is hard; it includes as a special case, for example, the general Boolean satisfaction problem.

Cardinality constrained convex optimization. As another broad special case of (1), consider the case $\mathcal{C} = \{z \in \mathbf{R}^q \mid \mathbf{card}(z) \leq k,\ \|z\|_\infty \leq M\}$, where $\mathbf{card}(z)$ is the number of nonzero elements of z, and k and M are given. We call this the general *cardinality-constrained convex problem*. It arises in many interesting applications, such as regressor selection.

Other special cases. As we will see in §6, many (hard) problems can be formulated in the form (1). More examples include regressor selection, 3-SAT, circle packing, the traveling salesman problem, factor analysis modeling, job selection, the maximum coverage problem, inexact graph isomorphism, and many more.

1.3 Convex relaxation

Convex relaxation of a set. A compact set $\mathcal{C}$ always has a tractable *convex relaxation*. By this we mean a (modest-sized) set of convex inequality and linear equality constraints that hold for every $z \in \mathcal{C}$:

$$z \in \mathcal{C} \implies h_i(z) \leq 0, \quad i = 1, \ldots, s, \qquad Fz = g.$$

We will assume that these relaxation constraints are included in the convex constraints of (1). Adding these relaxation constraints to the original problem yields an equivalent problem (since the added constraints are redundant), but can improve the convergence of any method, global or heuristic. By tractable, we mean that the number of added constraints is modest, and in particular, polynomial in q.

For example, when $\mathcal{C} = \{0, 1\}^q$, we have the inequalities $0 \leq z_i \leq 1$, $i = 1, \ldots, q$. (These inequalities define the convex hull of $\mathcal{C}$, *i.e.*, all other convex inequalities that hold for all $z \in \mathcal{C}$ are implied by them.) When

$$\mathcal{C} = \{z \in \mathbf{R}^q \mid \mathbf{card}(z) \leq k,\ \|z\|_\infty \leq M\},$$

we have the convex inequalities

$$\|z\|_1 \leq kM, \qquad \|z\|_\infty \leq M.$$

(These inequalities define the convex hull of $\mathcal{C}$.) For general compact $\mathcal{C}$ the inequality $\|z\|_\infty \leq M$ will always be a convex relaxation for some M.

Relaxed problem. If we remove the nonconvex constraint $z \in \mathcal{C}$, we get a *convex relaxation* of the original problem:

$$\begin{array}{ll} \text{minimize} & f_0(x,z) \\ \text{subject to} & f_i(x,z) \leq 0, \quad i = 1,\ldots,m \\ & Ax + Bz = c. \end{array} \tag{2}$$

(Recall that convex equalities and inequalities known to hold for $z \in \mathcal{C}$ have been incorporated in the convex constraints.) The relaxed problem is convex; its optimal value is a lower bound on the optimal value $p^\star$ of (1). A solution (x^*, z^*) to problem (2) need not satisfy $z^* \in \mathcal{C}$, but if it does, the pair (x^*, z^*) is optimal for (1).

1.4 Projections and approximate projections

Our methods will make use of tractable projection, or tractable approximate projection, onto the set $\mathcal{C}$. The usual Euclidean projection onto $\mathcal{C}$ will be denoted Π. (It need not be unique when $\mathcal{C}$ is not convex.) By approximate projection, we mean any function $\hat{\Pi} : \mathbf{R}^q \to \mathcal{C}$ that satisfies $\hat{\Pi}(z) = z$ for $z \in \mathcal{C}$. For example, when $\mathcal{C} = \{0,1\}^q$, exact projection is given by rounding the entries to $\{0,1\}$.

As a less trivial example, consider the cardinality-constrained problem. The projection of z onto $\mathcal{C}$ is given by

$$\left(\Pi\left(z\right)\right)_i = \begin{cases} M & z_i > M,\ i \in \mathcal{I} \\ -M & z_i < -M,\ i \in \mathcal{I} \\ z_i & |z_i| \leq M,\ i \in \mathcal{I} \\ 0 & i \notin \mathcal{I}, \end{cases}$$

where $\mathcal{I} \subseteq \{1,\ldots,q\}$ is a set of indices of k largest values of $|z_i|$. We will describe many projections, and some approximate projections, in §4.

1.5 Residual and merit functions

For any (x,z) with $z \in \mathcal{C}$, we define the *constraint residual* as

$$r(x,z) = \sum_{i=1}^{m} (f_i(x,z))_+ + \|Ax + Bz - c\|_1,$$

where $(u)_+ = \max\{u, 0\}$ denotes the positive part; (x, z) is feasible if and only if $r(x, z) = 0$. Note that $r(x, z)$ is a convex function of (x, z). We define the *merit function* of a pair (x, z) as

$$\eta(x, z) = f_0(x, z) + \lambda r(x, z),$$

where $\lambda > 0$ is a parameter. The merit function is also a convex function of (x, z).

When $\mathcal{C}$ is convex and the problem is feasible, minimizing $\eta(x, z)$ for large enough λ yields a solution of the original problem (1) (that is, the residual is a so-called exct penalty function); when the problem is not feasible, it tends to find approximate solutions that satisfy many of the constraints [HM79, PG89, Fle73].

We will use the merit function to judge candidate approximate solutions (x, z) with $z \in \mathcal{C}$; that is, we take a pair with lower merit function value to be a better approximate solution than one with higher merit function value. For some problems (for example, unconstrained problems) it is easy to find feasible points, so all candidate points will be feasible. The merit function then reduces to the objective value. At the other extreme, for feasibility problems the objective is zero, and goal is to find a feasible point. In this case the merit function reduces to $\lambda r(x, z)$, *i.e.*, a positive multiple of the residual function.

1.6 Solution methods

In this section we describe various methods for solving the problem (1), either exactly (globally) or approximately.

Global methods Depending on the set $\mathcal{C}$, the problem (1) can be solved globally by a variety of algorithms, including (or mixing) branch-and-bound [LW66, NF77, BJS94], branch-and-cut [PR91, TS05b, SM99], semidefinite hierarchies [SA90], or even direct enumeration when $\mathcal{C}$ is a finite set. In each iteration of these methods, a convex optimization problem derived from (1) is solved, with $\mathcal{C}$ removed, and (possibly) additional variables and convex constraints added. These global methods are generally thought to have high worst-case complexities and indeed can be very slow in practice, even for modest size problem instances.

Local solution methods and heuristics A local method for (1) solves a modest number of convex problems, in an attempt to find a good approximate solution, *i.e.*, a pair (x, z) with $z \in \mathcal{C}$ and a low value of the merit function $\eta(x, z)$. For a feasibility problem, we might hope to find a solution; and if not, find one with a small constraint residual. For a general problem, we can hope to find a feasible point with low objective value, ideally near the lower bound on $p^\star$ from the relaxed problem. If we cannot find any feasible points, we can settle for a pair (x, z) with $z \in \mathcal{C}$ and low merit function value. All of these methods are heuristics, in the sense that they cannot in general be guaranteed to find an optimal, or even good, or even feasible, point in only a modest number of iterations.

There are of course many heuristics for the general problem (1) and for many of its special cases. For example, any global optimization method can be stopped after some modest number of iterations; we take the best point found (in terms of the merit function)

as our approximate solution. We will discuss some local search methods, including neighbor search and polishing, in §2.

Existing solvers There are numerous open source and commercial solvers that can handle problems with nonconvex constraints. We only mention a few of them here. Gurobi [GO15], CPLEX [CPL09], MOSEK [ApS15] provide global methods for mixed integer linear programs, mixed integer quadratic programs, and mixed integer second order cone programs. BARON [TS05a], Couenne [CPL09], and SCIP [Ach09] use global methods for nonlinear programs and mixed integer nonlinear programs. Bonmin [BBC+08] and Knitro [BNW06] provide global methods for mixed integer convex programs and heuristic methods for mixed integer nonliner programs. IPOPT [WLMK09] and NLopt [Joh14] use heuristic methods for nonlinear programs.

1.7 Our approach

The purpose of this paper is to describe a general system for heuristic solution of (1), based on solving a *modest* number of convex problems derived from (1). By heuristic, we mean that the algorithm need not find an optimal point, or indeed, even a feasible point, even when one exists. We would hope that for many feasible problem instances from some application, the algorithm does find a feasible point, and one with objective not too far from the optimal value. The disadvantage of a heuristic over a global method is clear and simple: it need not find an optimal point. The advantage of a heuristic is that it can be (and often is) dramatically faster to carry out than a global method. Moreover there are many applications where a heuristic method for (1) is sufficient. This might be the case when the objective and constraints are already approximations of what we really want, so the added effort of solving it globally is not worth it.

ADMM. One of the heuristic methods described in this paper is based on the alternating directions method of multipliers (ADMM), an operator splitting algorithm originally devised to solve convex optimization problems [BPC+11]. We call this heuristic nonconvex alternating directions method of multipliers (NC-ADMM). The idea of using ADMM as a heuristic to solve nonconvex problems was mentioned in [BPC+11, Ch. 9], and has been explored by Yedidia and others [DBEY13] as a message passing algorithm. Consensus ADMM has been used for general quadratically constrained quadratic programming in [HS16]. In [XYWZ12], ADMM has been applied to non-negative matrix factorization with missing values. ADMM also has been used for real and complex polynomial optimization models in [JMZ14],for constrained tensor factorization in [LS14], and for optimal power flow in [Ers14]. ADMM is a generalization of the method of multipliers [Hes69, Ber14], and there is a long history of using the method of multipliers to (attempt to) solve nonconvex problems [Cha12, CW13, Hon14, HLR14, PCZ15, WXX14, LP15]. Several related methods, such as the Douglas-Rachford method [EB92] or Spingarn's method of partial inverses [Spi85], could just as well have been used.

Our contribution. The paper has the following structure. In §3 we discuss local search methods and describe how they can be used as solution improvement methods. This will enable us to study simple but sophisticated methods such as relax-round-polish, and iterative neighbor search. In §4 we catalog a variety of nonconvex sets for which Euclidean projection or approximated projection is easily evaluated and, when applicable, we discuss relaxations, restrictions, and the set of neighbors for a given point. In §5 we discuss an implementation of our general system for heuristic solution NCVX, as an extension of CVXPY [DCB14], a Python package for formulating and solving convex optimization problems. The object-oriented features of CVXPY make the extension particularly simple to implement. Finally, in §6 we demonstrate the performance of our methods on several example problems.

2 Local improvement methods

In this section we describe some simple general local search methods. These methods take a point $z \in \mathcal{C}$ and by performing a local search on z they find a candidate pair $(\hat{x}, \hat{z})$, with $\hat{z} \in \mathcal{C}$ and a lower merit function. We will see that for many applications these methods with a good initialization can be used to obtain an approximate solution. We will also see how we can use these methods to improve solution candidates from other heuristics, hence we refer to these methods as *solution improvement*.

2.1 Polishing

Convex restriction. We can have a tractable *convex restriction* of $\mathcal{C}$, that includes a given point in $\mathcal{C}$. This means that for each point $\tilde{z} \in \mathcal{C}$, we have a set of convex equalities and inequalities on z, that hold for $\tilde{z}$, and imply $z \in \mathcal{C}$. We denote the set of points that satisfy the restrictions as $\mathcal{C}^{\rm rstr}(\tilde{z})$, and call this set the *restriction* of $\mathcal{C}$ at $\tilde{z}$. The restriction set $\mathcal{C}^{\rm rstr}(\tilde{z})$ is convex, and satisfies $\tilde{z} \in \mathcal{C}^{\rm rstr}(\tilde{z}) \subseteq \mathcal{C}$. The trivial restriction is given by $\mathcal{C}^{\rm rstr}(\tilde{z}) = \{\tilde{z}\}$.

When $\mathcal{C}$ is discrete, for example $\mathcal{C} = \{0,1\}^q$, the trivial restriction is the only restriction. In other cases we can have interesting nontrivial restrictions, as we will see below. For example, with $\mathcal{C} = \{z \in \mathbf{R}^q \mid \mathbf{card}(z) \leq k,\ \|z\|_\infty \leq M\}$, we can take as restriction $\mathcal{C}^{\rm rstr}(\tilde{z})$, the set of vectors z with the same sparsity pattern as $\tilde{z}$, and $\|z\|_\infty \leq M$.

Polishing. Given any point $\tilde{z} \in \mathcal{C}$, we can replace the constraint $z \in \mathcal{C}$ with $z \in \mathcal{C}^{\rm rstr}(\tilde{z})$ to get the convex problem

$$\begin{array}{ll} \text{minimize} & \eta(x, z) \\ \text{subject to} & z \in \mathcal{C}^{\rm rstr}(\tilde{z}), \end{array} \tag{3}$$

with variables x, z. (When the restriction $\mathcal{C}^{\rm rstr}(\tilde{z})$ is the trivial one, *i.e.*, a singleton, this is equivalent to fixing $z = \tilde{z}$ and minimizing over x.) We call this problem the *convex restriction* of (1) at the point $\tilde{z}$. The restricted problem is convex, and its optimal value is an upper bound on $p^\star$.

As a simple example of polishing consider the mixed integer convex problem. The only restriction is the trivial one, so the polishing problem for a given Boolean vector $\tilde{z}$ simply fixes the values of the Boolean variables, and solves the convex problem over the remaining variables, *i.e.*, x. For the cardinality-constrained convex problem, polishing fixes the sparsity pattern of z and solves the resulting convex problem over z and x.

For problems with nontrivial restrictions, we can solve the polishing problem repeatedly until convergence. In other words we can use the output of the polishing problem as an initial point for another polishing problem and keep iterating until convergence or until a maximum number of iterations is reached. This technique is called *iterated polishing* and described in algorithm 1.

Algorithm 1 Iterated polishing

Input: $\tilde{z}$
do
    $z^{\mathrm{old}} \leftarrow \tilde{z}$.
    Find $(\tilde{x}, \tilde{z})$ by solving the polishing problem with restriction $z \in \mathcal{C}^{\mathrm{rstr}}(z^{\mathrm{old}})$.
while $\tilde{z} \neq z^{\mathrm{old}}$
return $(\hat{x}, \hat{z})$.

If there exists a point $\tilde{x}$ such that $(\tilde{x}, \tilde{z})$ is feasible, the restricted problem is feasible too. The restricted problem need not be feasible in general, but if it is, with solution $(\hat{x}, \hat{z})$, then the pair $(\hat{x}, \hat{z})$ is feasible for the original problem (1) and satisfies $f(\hat{x}, \hat{z}) \leq f(\tilde{x}, \tilde{z})$ for any $\tilde{x}$ for which $(\tilde{x}, \tilde{z})$ is feasible. So polishing can take a point $\tilde{z} \in \mathcal{C}$ (or a pair $(\tilde{x}, \tilde{z})$) and produce another pair $(\hat{x}, \hat{z})$ with a possibly better objective value.

2.2 Relax-round-polish

With the simple tools described so far (*i.e.*, relaxation, polishing, and projection) we can create several heuristics for approximately solving the problem (1). A basic version solves the relaxation, projects the relaxed value of z onto $\mathcal{C}$, and then polishes the result.

Algorithm 2 Relax-round-polish heuristic

Solve the convex relaxation (2) to obtain $(x^{\mathrm{rlx}}, z^{\mathrm{rlx}})$.
Find $z^{\mathrm{rnd}} = \Pi(z^{\mathrm{rlx}})$.
Find $(\hat{x}, \hat{z})$ by solving the polishing problem with restriction $z \in \mathcal{C}^{\mathrm{rstr}}(z^{\mathrm{rnd}})$.

Note that in the first step we also obtain a lower bound on the optimal value $p^\star$; in the polishing step we obtain an upper bound, and a feasible pair $(\hat{x}, \hat{z})$ that achieves the upper bound (provided that polishing is successful). The best outcome is for these bounds to be equal, which means that we have found a (global) solution of (1) (for this problem instance). But relax-round-polish can fail; for example, it can fail to find a feasible point even though one exists.

Many variations on relax-round-polish are possible. We can introduce randomization by replacing the round step with

$$z^{\mathrm{rnd}} = \Pi(z^{\mathrm{rlx}} + w),$$

where w is a random vector. We can repeat this heuristic with K different random instances of w. For each of K samples of w, we polish, giving us a set of K candidate approximate solutions. We then take as our final approximate solution the best among these K candidates, *i.e.*, the one with least merit function.

2.3 Neighbor search

Neighbors. We describe the concept of *neighbors* for a point $z \in \mathcal{C}$ when $\mathcal{C}$ is discrete. The set of neighbors of a point $z \in \mathcal{C}$, denoted $\mathcal{C}^{\mathrm{ngbr}}(z)$, is the set of points with distance one from z in a natural (integer valued) distance, which depends on the set $\mathcal{C}$. For example for the set of Boolean vectors in $\mathbf{R}^n$ we use *Hamming distance*, the number of entries in which two Boolean vectors differ. Hence neighbors of a Boolean vector z are the set of vectors that differ from z in one component. The distance between two permutation matrices is defined as the minimum number of swaps of adjacent rows and columns necessary to transform the first matrix into the second. With this distance, neighbors of a permutation matrix Z are the set of permutation matrices generated by swapping any two adjacent rows or columns in Z.

For Cartesian products of discrete sets we use the sum of distances. In this case, for $z = (z_1, z_2, \ldots, z_k) \in \mathcal{C} = \mathcal{C}_1 \times \mathcal{C}_2 \times \ldots \times \mathcal{C}_k$, neighbors of z are points of the form $(z_1, \ldots, z_{i-1}, \tilde{z}_i, z_{i+1}, \ldots, z_k)$ where $\tilde{z}_i$ is a neighbor of z_i in $\mathcal{C}_i$.

Basic neighbor search. We introduced polishing as a tool that can find a pair $(\hat{x}, \hat{z})$ given an input $\tilde{z} \in \mathcal{C}$ by solving a sequence of convex problems. In basic neighbor search we solve the polishing problem for $\tilde{z}$ and all neighbors of $\tilde{z}$ and return the pair (x^*, z^*) with the smallest merit function value. In practice, we can sample from $\mathcal{C}^{\mathrm{ngbr}}(\tilde{z})$ instead of iterating over all points in $\mathcal{C}^{\mathrm{ngbr}}(\tilde{z})$ if $\left|\mathcal{C}^{\mathrm{ngbr}}(\tilde{z})\right|$ is large.

Algorithm 3 Basic neighbor search

Input: $\tilde{z}$
Initialize $(x_{\mathrm{best}}, z_{\mathrm{best}}) = \emptyset$, $\eta_{\mathrm{best}} = \infty$.
for $\hat{z} \in \{\tilde{z}\} \cup \mathcal{C}^{\mathrm{ngbr}}(\tilde{z})$ **do**
    Find (x^*, z^*), by solving the polishing problem (3), with constraint $z \in \mathcal{C}^{\mathrm{rstr}}(\hat{z})$.
    if $\eta(x^*, z^*) < \eta_{\mathrm{best}}$ **then**
        $(x_{\mathrm{best}}, z_{\mathrm{best}}) = (x^*, z^*)$, $\eta_{\mathrm{best}} = \eta(x^*, z^*)$.
    end if
end for
return $(x_{\mathrm{best}}, z_{\mathrm{best}})$.

Iterated neighbor search. We can carry out the described neighbor search iteratively as follows. We maintain a current value of z, corresponding to the best pair (x, z) found so far. We then consider a neighbor of z and polish. If the new point (x, z) is better than the current best one, we reset our best and continue; otherwise we examine another neighbor. This is done until a maximum number of iterations is reached, or all neighbors of the current best z produce (under polishing) no better pairs. This procedure is sometimes called *hill climbing*, since it resembles an attempt to find the top of a mountain by repeatedly taking steps towards an ascent direction.

Algorithm 4 Iterative neighbor search

Input: $\tilde{z}$
Find $(x_{\rm best}, z_{\rm best})$ by solving the polishing problem (3)
$\eta_{\rm best} \leftarrow \eta(x_{\rm best}, z_{\rm best})$.
for $\hat{z} \in \mathcal{C}^{\rm ngbr}(z_{\rm best})$ **do**
    Find (x^*, z^*), by solving the polishing problem (3), with constraint $z \in \mathcal{C}^{\rm rstr}(\hat{z})$.
    if $\eta(x^*, z^*) < \eta_{\rm best}$ **then**
        $(x_{\rm best}, z_{\rm best}) = (x^*, z^*)$, $\eta_{\rm best} = \eta(x^*, z^*)$.
        Go to 4
    end if
end for
return $(x_{\rm best}, z_{\rm best})$.

Notice that when no neighbors are available for $\tilde{z} \in \mathcal{C}$, this algorithm reduces to simple polishing.

3 NC-ADMM

We already can use the simple tools described in the previous section as heuristics to find approximate solutions to problem (1). In this section, we describe the alternating direction method of multipliers (ADMM) as a mechanism to generate candidate points $\tilde{z}$ to carry out local search methods such as iterated neighbor search. We call this method nonconvex ADMM, or NC-ADMM.

3.1 ADMM

Define $\phi : \mathbf{R}^q \to \mathbf{R} \cup \{-\infty, +\infty\}$ such that $\phi(z)$ is the best objective value of problem (1) after fixing z. In other words,

$$\phi(z) = \inf_x \left\{ f_0(x, z) \mid f_i(x, z) \leq 0,\ i = 1, \ldots, m,\ Ax + Bz = c \right\}.$$

Notice that $\phi(z)$ can be $+\infty$ or $-\infty$ in case the problem is not feasible for this particular value of z, or problem (2) is unbounded below after fixing z. The function ϕ is convex, since

it is the partial minimization of a convex function over a convex set [BV04, §3.4.4]. It is defined over all points $z \in \mathbf{R}^q$, but we are interested in finding its minimum value over the nonconvex set $\mathcal{C}$. In other words, problem (1) can be formulated as

$$\begin{array}{ll} \text{minimize} & \phi(z) \\ \text{subject to} & z \in \mathcal{C}. \end{array} \tag{4}$$

As discussed in [BPC+11, Chapter 9], ADMM can be used as a heuristic to solve nonconvex constrained problems. ADMM has the form

$$\begin{aligned} w^{k+1} &:= \operatorname{argmin}_z \left(\phi(z) + (\rho/2)\|z - z^k + u^k\|_2^2\right) \\ z^{k+1} &:= \Pi\left(w^{k+1} - z^k + u^k\right) \\ u^{k+1} &:= u^k + w^{k+1} - z^{k+1}, \end{aligned} \tag{5}$$

where $\rho > 0$ is an algorithm parameter, k is the iteration counter, and Π denotes Euclidean projection onto $\mathcal{C}$ (which need not be unique when $\mathcal{C}$ is not convex).

The initial values u^0 and z^0 are additional algorithm parameters. We always set $u^0 = 0$ and draw z^0 randomly from a normal distribution $\mathcal{N}(0, \sigma^2 I)$, where $\sigma > 0$ is an algorithm parameter.

3.2 Algorithm subroutines

Convex proximal step Carrying out the first step of the algorithm, *i.e.*, evaluating the proximal operator of ϕ, involves solving the convex optimization problem

$$\begin{array}{ll} \text{minimize} & f_0(x, z) + (\rho/2)\|z - z^k + u^k\|_2^2 \\ \text{subject to} & f_i(x, z) \leq 0, \quad i = 1, \ldots, m, \\ & Ax + Bz = c, \end{array} \tag{6}$$

over the variables $x \in \mathbf{R}^n$ and $z \in \mathbf{R}^q$. This is the original problem (1), with the nonconvex constraint $z \in \mathcal{C}$ removed, and an additional convex quadratic term involving z added to the objective. We let (x^{k+1}, w^{k+1}) denote a solution of (6). If the problem (6) is infeasible, then so is the original problem (1); should this happen, we can terminate the algorithm with the certain conclusion that (1) is infeasible.

Projection The (nonconvex) projection step consists of finding the closest point in $\mathcal{C}$ to $w^{k+1} - z^k + u^k$. If more than one point has the smallest distance, we can choose one of the minimizers arbitrarily.

Dual update The iterate $u^k \in \mathbf{R}^q$ can be interpreted as a scaled dual variable, or as the running sum of the error values $w^{k+1} - z^k$.

3.3 Discussion

Convergence. When $\mathcal{C}$ is convex (and a solution of (1) exists), this algorithm is guaranteed to converge to a solution, in the sense that $f_0(x^{k+1}, w^{k+1})$ converges to the optimal value of the problem (1), and $w^{k+1} - z^{k+1} \to 0$, *i.e.*, $w^{k+1} \to \mathcal{C}$. See [BPC+11, §3] and the references therein for a more technical description and details. But in the general case, when $\mathcal{C}$ is not convex, the algorithm is not guaranteed to converge, and even when it does, it need not be to a global, or even local, minimum. Some recent progress has been made on understanding convergence in the nonconvex case [LP15].

Parameters. Another difference with the convex case is that the convergence and the quality of solution depends on ρ, whereas for convex problems this algorithm is guaranteed to converge to the optimal value regardless of the choice of ρ. In other words, in the convex case the choice of parameter ρ only affects the speed of the convergence, while in the nonconvex case the choice of ρ can have a critical role in the quality of approximate solution, as well as the speed at which this solution is found.

The optimal parameter selection for ADMM is still an active research area. In [GTSJ15] the optimal parameter selection for quadratic problems is discussed. In a more generalized setting, Giselsson discusses the optimal parameter selection for ADMM for strongly convex functions [GB14a, GB14b, GB14c]. The dependency of global and local convergence properties of ADMM on parameter choice has been studied in [HL12, Bol13].

Initialization. In the convex case the choice of initial point z^0 affects the number of iterations to find a solution, but not the quality of the solution. Unsurprisingly, the nonconvex case differs in that the choice of z^0 has a major effect on the the quality of the approximate solution. As with the choice of ρ, the initialization in the nonconvex case is currently an active area of research; see, *e.g.*, [HS16, LP15, TMBB15]. Getting the best possible results on a particular problem requires a careful and problem specific choice of initialization. We draw initial points randomly from $\mathcal{N}(0, \sigma^2 I)$ because we want a method that generalizes easily across many different problems.

3.4 Solution improvement

Now we describe two techniques to obtain better solutions after carrying out ADMM. The first technique relies on iterated neighbor search and the second one is using multiple restarts with random initial points in order to increase the chance of obtaining a better solution.

Iterated neighbor search After each iteration, we can carry out iterated polishing (as described in §2.3) with $\mathcal{C}^{\text{rstr}}(z^{k+1})$ to obtain $(\hat{x}^{k+1}, \hat{z}^{k+1})$. We will return the pair with the smallest merit function as the output of the algorithm.

Multiple restarts As we mentioned, we choose the initial value z^0 from a normal distribution $\mathcal{N}(0, \sigma^2 I)$. We can run the algorithm multiple times from different initial points to increase the chance of a feasible point with a smaller objective value.

3.5 Overall algorithm

The following is a summary of the algorithm with solution improvement.

Algorithm 5 NC-ADMM heuristic

Initialize $u^0 = 0$, $(x_{\text{best}}, z_{\text{best}}) = \emptyset$, $\eta_{\text{best}} = \infty$.
for algorithm repeats $1, 2, \ldots, M$ **do**
  Initialize $z^0 \sim \mathcal{N}(0, \sigma^2 I)$.
  for $k = 1, 2, \ldots, N$ **do**
    $(x^{k+1}, w^{k+1}) \leftarrow \operatorname{argmin}_z \left(\phi(z) + (\rho/2)\|z - z^k + u^k\|_2^2\right)$.
    $z^{k+1} \leftarrow \Pi\left(w^{k+1} - z^k + u^k\right)$.
    Use algorithm (4) on z^{k+1} to get the improved iterate $(\hat{x}, \hat{z})$.
    if $\eta(\hat{x}, \hat{z}) < \eta_{\text{best}}$ **then**
      $(x_{\text{best}}, z_{\text{best}}) \leftarrow (\hat{x}, \hat{z})$, $\eta_{\text{best}} = \eta(\hat{x}, \hat{z})$.
    end if
    $u^{k+1} \leftarrow u^k + w^{k+1} - z^{k+1}$.
  end for
end for
return $x_{\text{best}}, z_{\text{best}}$.

4 Projections onto nonconvex sets

In this section we catalog various nonconvex sets with their implied convex constraints which will be included in the convex constraints of problem (1). We also provide a Euclidean projection (or approximate projection) Π for these sets. Also, when applicable, we introduce a nontrivial restriction and set of neighbors.

4.1 Subsets of R

Booleans For $\mathcal{C} = \{0, 1\}$, a convex relaxation (in fact, the convex hull of $\mathcal{C}$) is $[0, 1]$. Also, a projection is simple rounding: $\Pi(z) = 0$ for $z \leq 1/2$, and $\Pi(z) = 1$ for $z > 1/2$. ($z = 1/2$ can be mapped to either point.) Moreover, $\mathcal{C}^{\text{ngbr}}(0) = \{1\}$ and $\mathcal{C}^{\text{ngbr}}(1) = \{0\}$.

Finite sets If $\mathcal{C}$ has M elements, the convex hull of $\mathcal{C}$ is the interval from the smallest to the largest element. We can project onto $\mathcal{C}$ with no more than $\log_2 M$ comparisons. For

each $z \in \mathcal{C}$ the set of neighbors of $\mathcal{C}$ are the immediate points to the right and left of z (if they exist).

Bounded integers Let $\mathcal{C} = \mathbf{Z} \cap [-M, M]$, where $M > 0$. The convex hull is the interval from the smallest to the largest element integer in $[-M, M]$, *i.e.*, $[-\lfloor M \rfloor, \lfloor M \rfloor]$. The projection onto $\mathcal{C}$ is simple: if $z > \lfloor M \rfloor$ ($z < -\lfloor M \rfloor$) then $\Pi(z) = \lfloor M \rfloor$ ($\Pi(z) = -\lfloor M \rfloor$). Otherwise, the projection of z can be found by simple rounding. For each $z \in \mathcal{C}$ the set of neighbors of $\mathcal{C}$ is $\{z-1, z+1\} \cap [-M, M]$

4.2 Subsets of $\mathbf{R}^n$

Boolean vectors with fixed cardinality Let $\mathcal{C} = \{z \in \{0,1\}^n \mid \mathbf{card}(z) = k\}$. Any $z \in \mathcal{C}$ satisfies $0 \leq z \leq 1$ and $\mathbf{1}^T z = k$. We can project $z \in \mathbf{R}^n$ onto $\mathcal{C}$ by setting the k entries of z with largest value to one and the remaining entries to zero. For any point $z \in \mathcal{C}$, the set of neighbors of z is all points generated by swapping an adjacent 1 and 0 in z.

Vectors with bounded cardinality Let $\mathcal{C} = \{x \in [-M, M]^n \mid \mathbf{card}(x) \leq k\}$, where $M > 0$ and $k \in \mathbf{Z}_+$. (Vectors $z \in \mathcal{C}$ are called k-sparse.) Any point $z \in \mathcal{C}$ satisfies $-M \leq z \leq M$ and $-Mk \leq \mathbf{1}^T z \leq Mk$. The projection $\Pi(z)$ is found as follows

$$(\Pi(z))_i = \begin{cases} M & z_i > M,\ i \in \mathcal{I} \\ -M & z_i < -M,\ i \in \mathcal{I} \\ z_i & |z_i| \leq M,\ i \in \mathcal{I} \\ 0 & i \notin \mathcal{I}, \end{cases}$$

where $\mathcal{I} \subseteq \{1, \ldots, n\}$ is a set of indices of k largest values of $|z_i|$.

A restriction of $\mathcal{C}$ at $z \in \mathcal{C}$ is the set of all points in $[-M, M]^n$ that have the same sparsity pattern as z. For any point $z \in \mathcal{C}$, the set of neighbors of z are all points $x \in \mathcal{C}$ whose sparsity pattern $\tilde{x} \in \{0,1\}^n$ is a neighbor of z's sparsity pattern $\tilde{z} \in \{0,1\}^n$. In other words, $\tilde{x}$ can be obtained by swapping an adjacent 1 and 0 in $\tilde{z}$.

Quadratic sets Let $\mathbf{S}^n_+$ and $\mathbf{S}^n_{++}$ denote the set of $n \times n$ symmetric positive semidefinite and symmetric positive definite matrices, respectively. Consider the set

$$\mathcal{C} = \{z \in \mathbf{R}^n \mid \alpha \leq z^T A z + 2b^T z \leq \beta\},$$

where $A \in \mathbf{S}^n_{++}$, $b \in \mathbf{R}^n$, and $\beta \geq \alpha \geq -b^T A^{-1} b$. We assume $\alpha \geq -b^T A^{-1} b$ because $z^T A z + 2b^T z \geq -b^T A^{-1} b$ for all $z \in \mathbf{R}^n$. Any point $z \in \mathcal{C}$ satisfies the convex inequality $z^T A z + 2b^T z \leq \beta$.

We can find the projection onto $\mathcal{C}$ as follows. If $z^T A z + 2b^T z > \beta$, it suffices to solve

$$\begin{array}{ll} \text{minimize} & \|x - z\|_2^2 \\ \text{subject to} & x^T A x + 2b^T x \leq \beta, \end{array} \tag{7}$$

and if $z^T Az + 2b^T z < \alpha$, it suffices to solve

$$\begin{array}{ll} \text{minimize} & \|x - z\|_2^2 \\ \text{subject to} & x^T Ax + 2b^T x \geq \alpha. \end{array} \tag{8}$$

(If $\alpha \leq z^T Az + 2b^T z \leq \beta$, clearly $\Pi(z) = z$.) The first problem is a convex quadratically constrained quadratic program and the second problem can be solved by solving a simple semidefinite program as described in [BV04, Appendix B]. Furthermore, there is a more efficient way to find the projection by finding the roots of a single-variable polynomial of degree $2p + 1$, where p is the number of distinct eigenvalues of A [HS16, Hma10]. Note that the projection can be easily found even if A is not positive definite; we assume $A \in \mathbf{S}^n_{++}$ only to make $\mathcal{C}$ compact and have a useful convex relaxation.

A restriction of $\mathcal{C}$ at $z \in \mathcal{C}$ is the set

$$\mathcal{C}^{\text{rstr}}(z) = \{x \in \mathbf{R}^n \mid \frac{x^T Az + b^T(x + z) + b^T A^{-1} b}{\sqrt{z^T Az + 2b^T z + b^T A^{-1} b}} \geq \sqrt{\alpha + b^T A^{-1} b}, \quad x^T Ax + 2b^T x \leq \beta\}.$$

Recall that $z^T Az + 2b^T z + b^T A^{-1} b \geq 0$ for all $z \in \mathbf{R}^n$ and we assume $\alpha \geq -b^T A^{-1} b$, so $\mathcal{C}^{\text{rstr}}(z)$ is always well defined.

Annulus and sphere Consider the set

$$\mathcal{C} = \{z \in \mathbf{R}^n \mid r \leq \|z\|_2 \leq R\},$$

where $R \geq r$.

Any point $z \in \mathcal{C}$ satisfies $\|z\|_2 \leq R$. We can project $z \in \mathbf{R}^n \setminus \{0\}$ onto $\mathcal{C}$ by the following scaling

$$\Pi(z) = \begin{cases} rz/\|z\|_2 & \text{if } \|z\|_2 < r \\ z & \text{if } z \in \mathcal{C} \\ Rz/\|z\|_2 & \text{if } \|z\|_2 > R, \end{cases}$$

If $z = 0$, any point with Euclidean norm r is a valid projection.

A restriction of $\mathcal{C}$ at $z \in \mathcal{C}$ is the set

$$\mathcal{C}^{\text{rstr}}(z) = \{x \in \mathbf{R}^n \mid x^T z \geq r\|z\|_2, \ \|x\|_2 \leq R\}.$$

Notice that if $r = R$, then $\mathcal{C}$ is a sphere and the restriction will be a singleton.

Box complement and cube surface Consider the set

$$\mathcal{C} = \{z \in \mathbf{R}^n \mid a \leq \|z\|_\infty \leq b\}.$$

Any point $z \in \mathcal{C}$ satisfies $\|z\|_\infty \leq b$. For any point z we can find the projection $\Pi(z)$ by projecting z component-wise onto $[a, b]$.

Given $z \in \mathcal{C}$ we can obtain a restriction by finding $k = \operatorname{argmin}_i \max\{|z_i|, a\}$ and if $z_k \geq 0$ then

$$\mathcal{C}^{\text{rstr}}(z) = \{x \mid x_k \geq a,\ \|x\|_\infty \leq b\}.$$

If $z_k < 0$, then

$$\mathcal{C}^{\text{rstr}}(z) = \{x \mid x_k \leq -a,\ \|x\|_\infty \leq b\}.$$

Notice that if $a = b$, then $\mathcal{C}$ is a cube surface.

4.3 Subsets of $\mathbf{R}^{m\times n}$

Remember that the projection of a point $X \in \mathbf{R}^{m\times n}$ on a set $\mathcal{C} \subset \mathbf{R}^{m\times n}$ is a point $Z \in \mathcal{C}$ such that the Frobenius norm $\|X - Z\|_{\text{F}}$ is minimized. As always, if there is more than one point Z that minimizes $\|X - Z\|_{\text{F}}$, we accept any of them.

Matrices with bounded singular values and orthogonal matrices Consider the set of $m \times n$ matrices whose singular values lie between 1 and α

$$\mathcal{C} = \{Z \in \mathbf{R}^{m\times n} \mid I \preceq Z^T Z \preceq \alpha^2 I\},$$

where $\alpha \geq 1$, and $A \preceq B$ means $B - A \in \mathbf{S}^n_+$. Any point $Z \in \mathcal{C}$ satisfies $\|Z\|_2 \leq \alpha$.

If $Z = U\Sigma V^T$ is the singular value decomposition of Z with singular values $(\sigma_z)_{\min\{m,n\}} \leq \cdots \leq (\sigma_z)_1$ and $X \in \mathcal{C}$ with singular values $(\sigma_x)_{\min\{m,n\}} \leq \cdots \leq (\sigma_x)_1$, according to the von Neumann trace inequality [Neu37] we will have

$$\mathbf{Tr}(Z^T X) \leq \sum_{i=1}^{\min\{m,n\}} (\sigma_z)_i (\sigma_x)_i.$$

Hence

$$\|Z - X\|_F^2 \geq \sum_{i=1}^{\min\{m,n\}} \left((\sigma_z)_i - (\sigma_x)_i\right)^2,$$

with equality when $X = U\,\mathbf{diag}(\sigma_x)V^T$. This inequality implies that $\Pi(Z) = U\tilde{\Sigma}V^T$, where $\tilde{\Sigma}$ is a diagonal matrix and $\tilde{\Sigma}_{ii}$ is the projection of Σ_{ii} on interval $[1, \alpha]$. When $Z = 0$, the projection $\Pi(Z)$ is any matrix.

Given $Z = U\Sigma V^T \in \mathcal{C}$, we can have the following restriction [BHA15]

$$\mathcal{C}^{\text{rstr}}(Z) = \{X \in \mathbf{R}^{m\times n} \mid \|X\|_2 \leq \alpha,\ V^T X^T U + U^T X V \succeq 2I\}.$$

(Notice that $X \in \mathcal{C}^{\text{rstr}}(Z)$ satisfies $X^T X \succeq I + (X - UV^T)^T(X - UV^T) \succeq I$.)

There are several noteworthy special cases. When $\alpha = 1$ and $m = n$ we have the set of orthogonal matrices. In this case, the restriction will be a singleton. When $n = 1$, the set $\mathcal{C}$ is equivalent to the annulus $\{z \in \mathbf{R}^m \mid 1 \leq \|z\|_2 \leq \alpha\}$.

Matrices with bounded rank Let $\mathcal{C} = \{Z \in \mathbf{R}^{m\times n} \mid \mathbf{Rank}(Z) \leq k,\ \|Z\|_2 \leq M\}$. Any point $Z \in \mathcal{C}$ satisfies $\|Z\|_2 \leq M$ and $\|Z\|_* \leq Mk$, where $\|\cdot\|_*$ denotes the trace norm. If $Z = U\Sigma V^T$ is the singular value decomposition of Z, we will have $\Pi(Z) = U\tilde{\Sigma}V^T$, where $\tilde{\Sigma}$ is a diagonal matrix with $\tilde{\Sigma}_{ii} = \min\{\Sigma_{ii}, M\}$ for $i = 1, \ldots k$, and $\tilde{\Sigma}_{ii} = 0$ otherwise.

Given a point $Z \in \mathcal{C}$, we can write the singular value decomposition of Z as $Z = U\Sigma V^T$ with $U \in \mathbf{R}^{m\times k}$, $\Sigma \in \mathbf{R}^{r\times r}$ and $V \in \mathbf{R}^{n\times k}$. A restriction of $\mathcal{C}$ at Z is

$$\mathcal{C}^{\mathrm{rstr}}(Z) = \{U\tilde{\Sigma}V^T \mid \tilde{\Sigma} \in \mathbf{R}^{r\times r}\}.$$

Assignment and permutation matrices The set of *assignment matrices* are Boolean matrices with exactly one non-zero element in each column and at most one non-zero element in each row. (They represent an assignment of the columns to the rows.) In other words, the set of assignment matrices on $\{0,1\}^{m\times n}$, where $m \geq n$, satisfy

$$\begin{array}{ll} \sum_{j=1}^{n} Z_{ij} \leq 1, & i = 1, \ldots, m \\ \sum_{i=1}^{m} Z_{ij} = 1, & j = 1, \ldots, n. \end{array}$$

These two sets of inequalities, along with $0 \leq Z_{ij} \leq 1$ are the implied convex inequalities. When $m = n$, this set becomes the set of permutation matrices, which we show by $\mathcal{P}_n$.

Projecting $Z \in \mathbf{R}^{m\times n}$ (with $m \geq n$) onto the set of assignment matrices involves choosing an entry from each column of Z such that no two chosen entries are from the same row and the sum of chosen entries is maximized. Assuming that the entries of Z are the weights of edges in a bipartite graph, the projection onto the set of assignment matrices will be equivalent to finding a maximum-weight matching in a bipartite graph. The Hungarian method [Kuh05] is a well-know polynomial time algorithm to find the maximum weight matching, and hence also the projection onto assignment matrices.

The neighbors of an assignment or permutation matrix $Z \in \mathbf{R}^{m\times n}$ are the matrices generated by swapping two adjacent rows or columns of Z.

Hamiltonian cycles A *Hamiltonian cycle* is a cycle in a graph that visits every node exactly once. Every Hamiltonian cycle in a complete graph can be represented by its adjacency matrix, for example

$$\begin{bmatrix} 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \end{bmatrix}$$

represents a Hamiltonian cycle that visits nodes $(3, 2, 4, 1)$ sequentially. Let $\mathcal{H}_n$ be the set of $n \times n$ matrices that represent a Hamiltonian cycle.

Every point $Z \in \mathcal{H}_n$ satisfies $0 \leq Z_{ij} \leq 1$ for $i, j = 1, \ldots, n$, and $Z = Z^T$, $(1/2)Z\mathbf{1} = \mathbf{1}$, and

$$2\mathbf{I} - Z + 4\frac{\mathbf{1}\mathbf{1}^T}{n} \geq 2(1 - \cos\frac{2\pi}{n})\mathbf{I},$$

where $\mathbf{I}$ denotes the identity matrix. In order to see why the last inequality holds, it's enough to notice that $2\mathbf{I} - Z$ is the Laplacian of the cycle represented by Z [Mer94, AM85]. It can

be shown that the smallest eigenvalue of $2\mathbf{I} - Z$ is zero (which corresponds to the eigenvector $\mathbf{1}$), and the second smallest eigenvalue of $2\mathbf{I} - Z$ is $2(1 - \cos\frac{2\pi}{n})$. Hence all eigenvalues of $2\mathbf{I} - Z + 4\frac{\mathbf{1}\mathbf{1}^T}{n}$ must be no smaller than $2(1 - \cos\frac{2\pi}{n})$.

We are not aware of a polynomial time algorithm to find the projection of a given real $n \times n$ matrix onto $\mathcal{H}_n$. We can find an *approximate projection* of Z by the following greedy algorithm: construct a graph with n vertices where the edge between i and j is weighted by z_{ij}. Start with the edge with largest weight and at each step, among all the edges that don't create a cycle, choose the edge with the largest weight (except for the last step where a cycle is created).

For a matrix $Z \in \mathcal{H}_n$, the set of neighbors of Z are matrices obtained after swapping two adjacent nodes, *i.e.*, matrices in form $P_{(i,j)}ZP_{(i,j)}^T$ where $Z_{ij} = 1$ and $P_{(i,j)}$ is a permutation matrix that transposes connected nodes i and j and keeps other nodes unchanged.

4.4 Combinations of sets

Cartesian product. Let $\mathcal{C} = \mathcal{C}_1 \times \cdots \times \mathcal{C}_k \subset \mathbf{R}^n$, where $\mathcal{C}_1, \ldots, \mathcal{C}_k$ are compact sets with known projections (or approximate projections). A convex relaxation of $\mathcal{C}$ is the Cartesian product $\mathcal{C}_1^{\text{rlx}} \times \cdots \times \mathcal{C}_k^{\text{rlx}}$, where $\mathcal{C}_i^{\text{rlx}}$ is the set described by the convex relaxation of $\mathcal{C}_i$. The projection of $z \in \mathbf{R}^n$ onto $\mathcal{C}$ is $(\Pi_1(z_1), \ldots, \Pi_k(z_k))$, where Π_i denotes the projection onto $\mathcal{C}_i$ for $i = 1, \ldots, k$.

A restriction of $\mathcal{C}$ at a point $z = (z_1, z_2, \ldots, z_k) \in \mathcal{C}$ is the Cartesian product $\mathcal{C}^{\text{rstr}}(z) = \mathcal{C}_1^{\text{rstr}}(z_1) \times \cdots \times \mathcal{C}_k^{\text{rstr}}(z_k)$. The neighbors of z are all points $(z_1, \ldots, z_{i-1}, \tilde{z}_i, z_{i+1}, \ldots, z_k)$ where $\tilde{z}_i$ is a neighbor of z_i in $\mathcal{C}_i$.

Union. Let $\mathcal{C} = \cup_{i=1}^k \mathcal{C}_i$, where $\mathcal{C}_1, \ldots, \mathcal{C}_k$ are compact sets with known projections (or approximate projections). A convex relaxation of $\mathcal{C}$ is the constraints

$$\begin{array}{ll} x_i \in \mathcal{C}_i^{\text{rlx}}, & i = 1, \ldots, k \\ s_i \in [0, 1], & i = 1, \ldots, k \\ z = \sum_{i=1}^k x_i & \\ \sum_{i=1}^k s_i = 1 & \\ \|x_i\|_\infty \leq M_i s_i, & i = 1, \ldots, k, \end{array}$$

where $\mathcal{C}_i^{\text{rlx}}$ is the set described by the convex relaxation of $\mathcal{C}_i$ and $M_i > 0$ is the minimum value such that $\|z_i\|_\infty \leq M_i$ holds for all $z_i \in \mathcal{C}_i$.

We can project $z \in \mathbf{R}^n$ onto $\mathcal{C}$ by projecting onto each set separately and keeping the projection closest to z:

$$\Pi(z) = \text{argmin}_{x \in \{\Pi_1(z), \cdots, \Pi_k(z)\}} \|z - x\|_2.$$

Here Π_i denotes the projection onto $\mathcal{C}_i$.

A restriction of $\mathcal{C}$ at a point z is $\mathcal{C}_i^{\text{rstr}}(z)$ for any $\mathcal{C}_i$ containing z. The neighbors of z are similarly the neighbors for any $\mathcal{C}_i$ containing z.

5 Implementation

We have implemented the NCVX Python package for modeling problems of the form (1) and applying the NC-ADMM heuristic, along with the relax-round-polish and relax methods. The NCVX package is an extension of CVXPY [DCB14]. The problem objective and convex constraints are expressed using standard CVXPY semantics. Nonconvex constraints are expressed implicitly by creating a variable constrained to lie in one of the sets described in §4. For example, the code snippet

```
x = Boolean()
```

creates a variable $x \in \mathbf{R}$ with the implicit nonconvex constraint $x \in \{0, 1\}$. The convex relaxation, in this case $x \in [0, 1]$, is also implicit in the variable definition. The source code for NCVX is available at `https://github.com/cvxgrp/ncvx`.

5.1 Variable constructors

The NCVX package provides the following functions for creating variables with implicit nonconvex constraints, along with many others not listed:

- `Boolean(n)` creates a variable $x \in \mathbf{R}^n$ with the implicit constraint $x \in \{0, 1\}^n$.
- `Integer(n, M)` creates a variable $x \in \mathbf{R}^n$ with the implicit constraints $x \in \mathbf{Z}^n$ and $\|x\|_\infty \leq \lfloor M \rfloor$.
- `Card(n, k, M)` creates a variable $x \in \mathbf{R}^n$ with the implicit constraints that at most k entries are nonzero and $\|x\|_\infty \leq M$.
- `Choose(n, k)` creates a variable $x \in \mathbf{R}^n$ with the implicit constraints that $x \in \{0, 1\}^n$ and has exactly k nonzero entries.
- `Rank(m, n, k, M)` creates a variable $X \in \mathbf{R}^{m \times n}$ with the implicit constraints $\mathbf{Rank}(X) \leq k$ and $\|X\|_2 \leq M$.
- `Assign(m, n)` creates a variable $X \in \mathbf{R}^{m \times n}$ with the implicit constraint that X is an assignment matrix.
- `Permute(n)` creates a variable $X \in \mathbf{R}^{n \times n}$ with the implicit constraint that X is a permutation matrix.
- `Cycle(n)` creates a variable $X \in \mathbf{R}^{n \times n}$ with the implicit constraint that X is the adjacency matrix of a Hamiltonian cycle.
- `Annulus(n,r,R)` creates a variable $x \in \mathbf{R}^n$ with the implicit constraint that $r \leq \|x\|_2 \leq R$.
- `Sphere(n, r)` creates a variable $x \in \mathbf{R}^n$ with the implicit constraint that $\|x\|_2 = r$.

5.2 Variable methods

Additionally, each variable created by the functions in §5.1 supports the following methods:

- `variable.relax()` returns a list of convex constraints that represent a convex relaxation of the nonconvex set $\mathcal{C}$, to which the variable belongs.
- `variable.project(z)` returns the Euclidean (or approximate) projection of z onto the nonconvex set $\mathcal{C}$, to which the variable belongs.
- `variable.restrict(z)` returns a list of convex constraints describing the convex restriction $\mathcal{C}^{\mathrm{rstr}}(z)$ at z of the nonconvex set $\mathcal{C}$, to which the variable belongs.
- `variable.neighbors(z)` returns a list of neighbors $\mathcal{C}^{\mathrm{ngbr}}(z)$ of z contained in the nonconvex set $\mathcal{C}$, to which the variable belongs.

Users can add support for additional nonconvex sets by implementing functions that return variables with these four methods.

5.3 Constructing and solving problems

To construct a problem of the form (1), the user creates variables $z_1, \ldots, z_k$ with the implicit constraints $z_1 \in \mathcal{C}_1, \ldots, z_k \in \mathcal{C}_k$, where $\mathcal{C}_1, \ldots, \mathcal{C}_k$ are nonconvex sets, using the functions described in §5.1. The variable z in problem (1) corresponds to the vector $(z_1, \ldots, z_k)$. The components of the variable x, the objective, and the constraints are constructed using standard CVXPY syntax.

Once the user has constructed a problem object, they can apply the following solve methods:

- `problem.solve(method="relax")` solves the convex relaxation of the problem.
- `problem.solve(method="relax-round-polish")` applies the relax-round-polish heuristic. Additional arguments can be used to specify the parameters K and λ. By default the parameter values are $K = 1$ and $\lambda = 10^4$. When $K > 1$, the first sample $w_1 \in \mathbf{R}^q$ is always 0. Subsequent samples are drawn i.i.d. from $N(0, \sigma^2 I)$, where σ is another parameter the user can set.
- `problem.solve(method="nc-admm")` applies the NC-ADMM heuristic. Additional arguments can be used to specify the number of starting points, the number of iterations the algorithm is run from each starting point, and the values of the parameters ρ, σ, and λ. By default the algorithm is run from 5 starting points for 50 iterations, the value of ρ is drawn uniformly from $[0, 1]$, and the other parameter values are $\sigma = 1$ and $\lambda = 10^4$. The first starting point is always $z^0 = 0$ and subsequent starting points are drawn i.i.d. from $\mathcal{N}(0, \sigma^2 I)$.

The relax-round-polish and NC-ADMM methods record the best point found $(x_{\text{best}}, z_{\text{best}})$ according to the merit function. The methods return the objective value $f_0(x_{\text{best}}, z_{\text{best}})$ and the residual $r(x_{\text{best}}, z_{\text{best}})$, and set the `value` field of each variable to the appropriate segment of x_{best} and z_{best}.

For example, consider the *regressor selection* problem, which we will discuss in §6.1. This problem can be formulated as

$$\begin{array}{ll} \text{minimize} & \|Ax-b\|_2^2 \\ \text{subject to} & \|x\|_\infty \leq M \\ & \mathbf{card}(x) \leq k, \end{array} \tag{9}$$

with decision variable $x \in \mathbf{R}^n$ and problem data $A \in \mathbf{R}^{m\times n}$, $b \in \mathbf{R}^m$, $M > 0$, and $k \in \mathbf{Z}_+$. The following code attempts to approximately solve this problem using our heuristic.

```
x = Card(n,k,M)
prob = Problem(Minimize(sum_squares(A*x-b)))
objective, residual = prob.solve(method="nc-admm")
```

The first line constructs a variable $x \in \mathbf{R}^n$ with the implicit constraints that at most k entries are nonzero, $\|x\|_\infty \leq M$, and $\|x\|_1 \leq kM$. The second line creates a minimization problem with objective $\|Ax-b\|_2^2$ and no constraints. The last line applies the NC-ADMM heuristic to the problem and returns the objective value and residual of the best point found.

6 Examples

In this section we apply the NC-ADMM heuristic to a wide variety of hard problems, *i.e.*, that generally cannot be solved in polynomial time. Extensive research has been done on specialized algorithms for each of the problems discussed in this section. Our intention is not to seek better performance than these specialized algorithms, but rather to show that our general purpose heuristic can yield decent results with minimal tuning. Unless otherwise specified, the algorithm parameters are the defaults described in §5. Whenever possible, we compare our heuristic to GUROBI [GO15], a commercial global optimization solver. Since our implementation of NC-ADMM supports minimal parallelization, we compare the number of convex subproblems solved (and not the solve time).

6.1 Regressor selection

We consider the problem of approximating a vector b with a linear combination of at most k columns of A with bounded coefficients. This problem can be formulated as

$$\begin{array}{ll} \text{minimize} & \|Ax-b\|_2^2 \\ \text{subject to} & \mathbf{card}(x) \leq k, \quad \|x\|_\infty \leq M, \end{array} \tag{10}$$

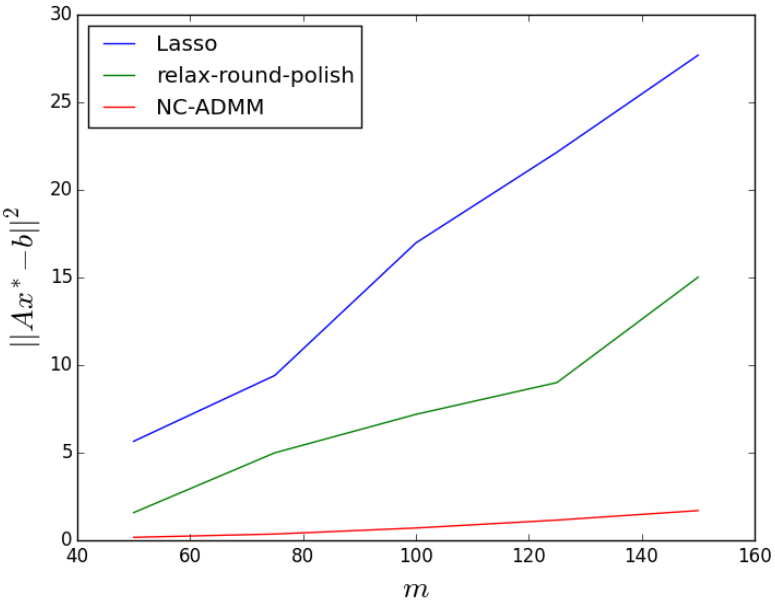

Figure 1: The average error of solutions found by Lasso, relax-round-polish, and NC-ADMM for 40 instances of the regressor selection problem.

with decision variable $x \in \mathbf{R}^n$ and problem data $A \in \mathbf{R}^{m\times n}$, $b \in \mathbf{R}^m$, $k \in \mathbf{Z}_+$, and $M > 0$. Lasso (least absolute shrinkage and selection operator) is a well-known heuristic for solving this problem by adding ℓ_1 regularization and minimizing $\|Ax - b\|_2^2 + \lambda\|x\|_1$. The value of λ is chosen as the smallest value possible such that $\mathbf{card}(x) \leq k$. (See [FHT01, §3.4] and [BV04, §6.3].)

Problem instances. We generated the matrix $A \in \mathbf{R}^{m\times 2m}$ with i.i.d. $\mathcal{N}(0,1)$ entries, and chose $b = A\hat{x} + v$, where $\hat{x}$ was drawn uniformly at random from the set of vectors satisfying $\mathbf{card}(\hat{x}) \leq \lfloor m/5 \rfloor$ and $\|x\|_\infty \leq 1$, and $v \in \mathbf{R}^m$ was a noise vector drawn from $\mathcal{N}(0, \sigma^2 I)$. We set $\sigma^2 = \|A\hat{x}\|^2/(400m)$ so that the signal-to-noise ratio was near 20.

Results. For each value of m, we generated 40 instances of the problem as described in the previous paragraph. Figure 1 compares the average sum of squares error for the x^* values found by the Lasso heuristic, relax-round-polish, and NC-ADMM. For Lasso, we solved the problem for 100 values of λ and then solved the polishing problem after fixing the sparsity pattern suggested by Lasso.

6.2 3-satisfiability

Given Boolean variables $x_1, \cdots, x_n$, a *literal* is either a variable or the negation of a variable, for example x_1 and $\neg x_2$. A *clause* is disjunction of literals (or a single literal), for example

$(\neg x_1 \vee x_2 \vee \neg x_3)$. Finally a formula is in conjunctive normal form (CNF) if it is a conduction of clauses (or a single clause), for example $(\neg x_1 \vee x_2 \vee \neg x_3) \wedge (x_1 \vee \neg x_2)$. Determining the satisfiability of a formula in conjunctive normal form where each clause is limited to at most three literals is called *3-satisfiability* or simply the *3-SAT* problem. It is known that 3-SAT is NP-complete, hence we do not expect to be able to solve a 3-SAT in general using our heuristic. A 3-SAT problem can be formulated as the following

$$\begin{array}{ll} \text{minimize} & 0 \\ \text{subject to} & Az \leq b, \\ & z \in \{0,1\}^n, \end{array} \tag{11}$$

where entries of A are given by

$$a_{ij} = \begin{cases} -1 & \text{if clause } i \text{ contains } x_j \\ 1 & \text{if clause } i \text{ contains } \neg x_j \\ 0 & \text{otherwise,} \end{cases}$$

and the entries of b are given by

$$b_i = (\text{number of negated literals in clause } i) - 1.$$

Problem instances. We generated 3-SAT problems with varying numbers of clauses and variables randomly as in [MSL92, LB14]. As discussed in [CA96], there is a threshold around 4.25 clauses per variable when problems transition from being feasible to being infeasible. Problems near this threshold are generally found to be hard satisfiability problems. We generated 10 instances for each choice of number of clauses and variables, verifying that each instance is feasible using GUROBI [GO15].

Results. We ran NC-ADMM heuristic on each instance, with 10 restarts, and 100 iterations, and we chose the step size $\rho = 10$. Figure 2 shows the fraction of instances solved correctly with NC-ADMM. We see that using this heuristic, satisfying assignments can be found consistently for up to 3.2 constraints per variable, at which point success starts to decrease. Problems in the gray region in figure 2 were not tested since they are infeasible with high probability. We also tried the relax-round-polish heuristic, but it often failed to solve problems with more than 50 clauses.

6.3 Circle packing

In *circle packing problem* we are interested in finding the smallest square in which we can place n non-overlapping circles with radii $r_1, \ldots, r_n$ [Gol70]. This problem has been studied extensively [Ste05, CKP08, CS03] and a database of densest known packings for different

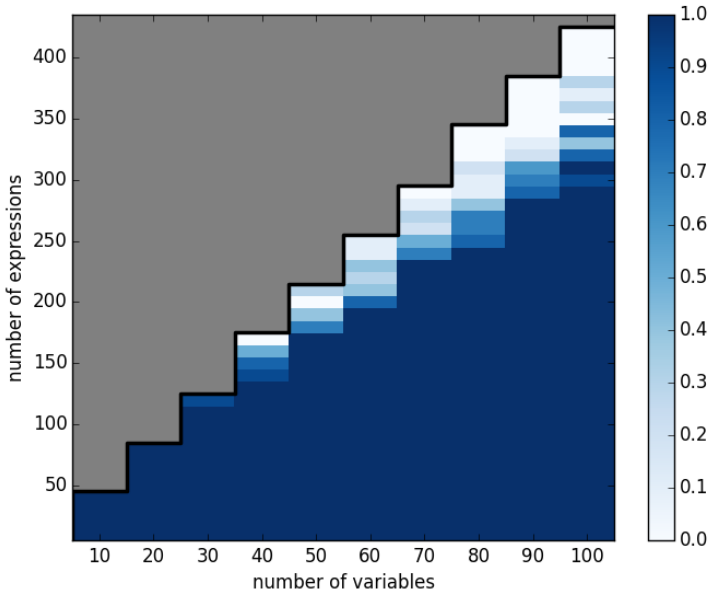

Figure 2: Fraction of runs for which a satisfying assignment to random 3-SAT problems were found for problems of varying sizes. The problems in the gray region were not tested.

numbers of circles can be found in [Spe13]. The problem can be formulated as

$$\begin{array}{ll} \text{minimize} & l \\ \text{subject to} & r_i\mathbf{1} \leq x_i \leq (l-r_i)\mathbf{1}, \quad i=1,\ldots,n \\ & x_i - x_j = z_{ij}, \quad i=1,\ldots,n-1, \quad j=i+1,\ldots,n \\ & 2\sum_{k=1}^{n} r_i \geq \|z_{ij}\|_2 \geq r_i + r_j, \quad i=1,\ldots,n-1, \quad j=i+1,\ldots,n, \end{array} \tag{12}$$

where $x_1, \ldots, x_n \in \mathbf{R}^2$ are variables representing the circle centers and $z_{12}, z_{13}, \ldots, z_{n-1,n} \in \mathbf{R}^2$ are additional variables representing the offset between pairs (x_i, x_j). Note that each z_{ij} is an element of an annulus.

Problem instances. We generated problems with different numbers of circles. Here we report the performance of the relax-round-polish heuristic for a problem with $n = 41$, in two cases: a problem with all circle radii equal to 0.5, and a problem where the radii were chosen uniformly at random from the interval $[0.2, 0.5]$.

Results. We run the relax-round-polish heuristic in both cases. For this problem, the heuristic is effectively equivalent to many well-known methods like the convex-concave procedure and the majorization-minimization (MM) algorithm. Figure 3 shows the packing found by our heuristic for $n = 41$. The obtained packing covers 78.68% of the area of the bounding square, which is close to the densest known packing, which covers 79.27% of the

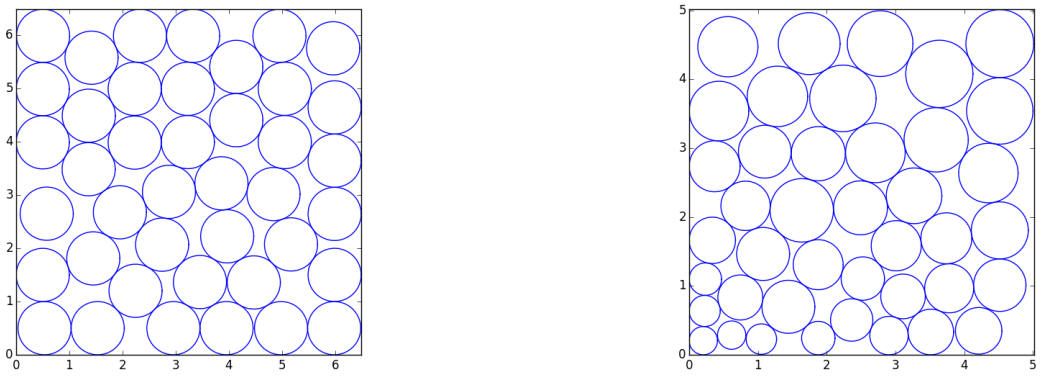

Figure 3: Packing for $n = 41$ circles with equal and different radii.

area. We observed that NC-ADMM is no more effective than relax-round-polish for this problem.

6.4 Traveling salesman problem

In the traveling salesman problem (TSP), we wish to find the minimum weight Hamiltonian cycle in a weighted graph. A Hamiltonian cycle is a path that starts and ends on the same vertex and visits each other vertex in the graph exactly once. Let G be a graph with n vertices and $D \in \mathbf{S}^n$ be the (weighted) adjacency matrix, *i.e.*, the real number d_{ij} denotes the distance between i and j. We can formulate the TSP problem for G as follows

$$\begin{array}{ll} \text{minimize} & (1/2)\,\mathbf{Tr}(D^T Z) \\ \text{subject to} & Z \in \mathcal{H}_n, \end{array} \tag{13}$$

where Z is the decision variable [Law85, Kru56, DFJ54, HPR13].

Problem instances. We generated $n = 75$ points in $[-1, 1]^2$. We set d_{ij} to be the Euclidean distance between points i and j.

Results. Figure 4 compares the Hamiltonian cycle found by the NC-ADMM heuristic, which had cost 14.47, with the optimal Hamiltonian cycle, which had cost 14.16. The cycle found by our heuristic has a few clearly suboptimal paths, but overall is a reasonable approximate solution. We ran NC-ADMM with 5 restarts and 100 iterations. GUROBI solved 4190 subproblems before finding a solution as good as that found by NC-ADMM, which solved only 500 subproblems. The relax-round-polish heuristic does not perform well on this problem. The best objective value found by the heuristic is 35.6.

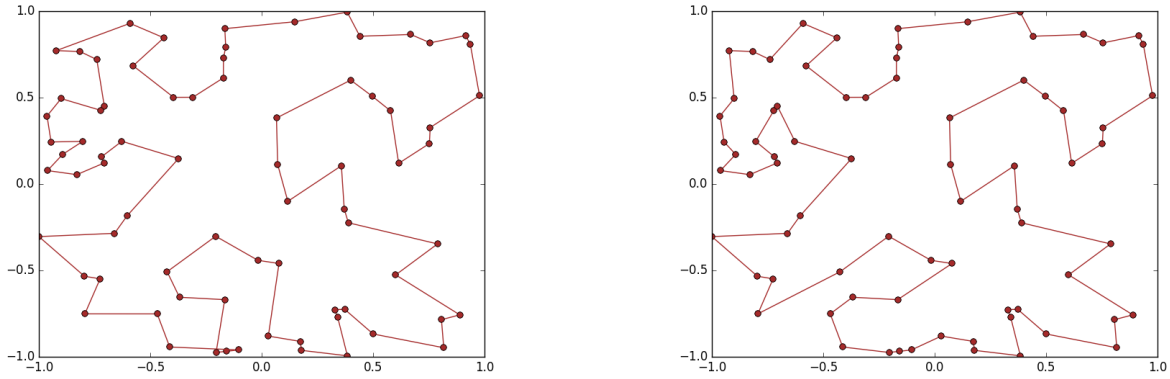

Figure 4: Left: Hamiltonian cycle found by NC-ADMM. Right: optimal Hamiltonian cycle, found using GUROBI.

6.5 Factor analysis model

The factor analysis problem decomposes a matrix as a sum of a low-rank and a diagonal matrix and has been studied extensively (for example in [SCPW12, NTGTB15]). It is also known as the *Frisch* scheme in the system identification literature [Kal85, DM93]. The problem is the following

$$\begin{array}{ll} \text{minimize} & \|\Sigma - \Sigma^{\text{lr}} - D\|_F^2 \\ \text{subject to} & D = \mathbf{diag}(d), \quad d \geq 0 \\ & \Sigma^{\text{lr}} \succeq 0 \\ & \mathbf{Rank}(\Sigma^{\text{lr}}) \leq k, \end{array} \tag{14}$$

where $\Sigma^{\text{lr}} \in \mathbf{S}_+^n$ and diagonal matrix $D \in \mathbf{R}^{n\times n}$ with nonnegative diagonal entries are the decision variables, and $\Sigma \in \mathbf{S}_+^n$ and $k \in \mathbf{Z}_+$ are problem data. One well-known heuristic for solving this problem is adding $\|\cdot\|_*$, or nuclear norm, regularization and minimizing $\|\Sigma - \Sigma^{\text{lr}} - D\|_F^2 + \lambda\|\Sigma^{\text{lr}}\|_*$. The value of λ is chosen as the smallest value possible such that $\mathbf{Rank}(\Sigma^{\text{lr}}) \leq k$. Since Σ^{lr} is positive semidefinite, $\|\Sigma^{\text{lr}}\|_* = \mathbf{Tr}(\Sigma^{\text{lr}})$.

Problem instances. We set $k = \lfloor n/2 \rfloor$ and generated the matrix $F \in \mathbf{R}^{n\times k}$ by drawing the entries i.i.d. from a standard normal distribution. We generated a diagonal matrix $\hat{D}$ with diagonal entries drawn i.i.d. from an exponential distribution with mean 1. We set $\Sigma = FF^T + \hat{D} + V$, where $V \in \mathbf{R}^{n\times n}$ is a noise matrix with entries drawn i.i.d. from $\mathcal{N}(0, \sigma^2)$. We set $\sigma^2 = \|FF^T + \mathbf{diag}(d)\|_F^2/(400n^2)$ so that the signal-to-noise ratio was near 20.

Results. Figure 5 compares the average sum of squares error for the Σ^{lr} and d values found by NC-ADMM, relax-round-polish, and the nuclear norm heuristic over 50 instances

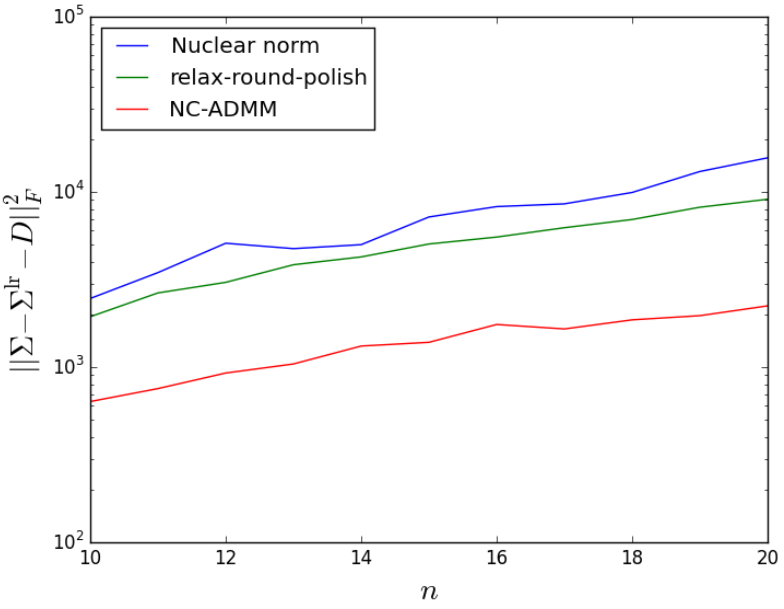

Figure 5: The average error of solutions found by the nuclear norm, relax-round-polish, and NC-ADMM heuristics for 50 instances of the factor analysis problem.

per value of n. We observe that the sum of squares error obtained by NC-ADMM is smaller than that obtained by the nuclear norm and relax-round-polish heuristics.

6.6 Job selection

In the job selection problem there are n jobs and m resources. Each job i consumes $A_{ij} \geq 0$ units of resource j, and up to $d_i > 0$ instances of job i can be accepted. Executing job i produces $c_i > 0$ units of profit. The goal is to maximize profit subject to the constraint that at most $b_j > 0$ units of each resource j are consumed. The job selection problem can be formulated as

$$\begin{array}{ll} \text{maximize} & c^T z \\ x\text{subject to} & Az \leq b \\ & 0 \leq z \leq d \\ & z \in \mathbf{Z}^n, \end{array} \tag{15}$$

where z is the decision variable and $A \in \mathbf{R}^{m \times n}$, $b \in \mathbf{R}^m_+$, and $d \in \mathbf{Z}^n_+$ are problem data. This problem is NP-hard in general. When $m = 1$, this problem is equivalent to the *knapsack problem*, which has been studied extensively; see, *e.g.*, [CB98, CK05].

Problem instances. We set $m = \lfloor n/10 \rfloor$ and generated $A \in \mathbf{R}^{m \times n}$ by randomly selecting $\lfloor mn/10 \rfloor$ entries to be nonzero. The nonzero entries were drawn i.i.d. from the uniform

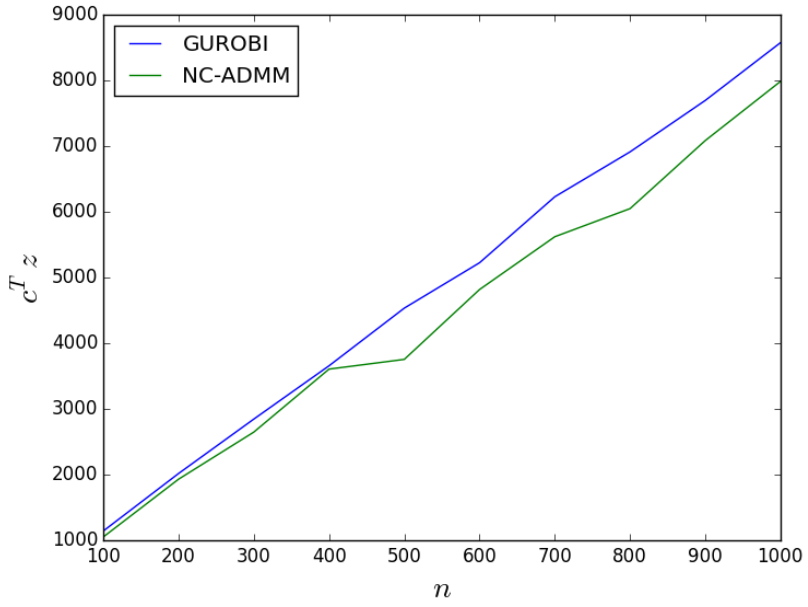

Figure 6: The average objective value for the job selection problem over 10 different instances.

distribution over $[0, 5]$. Entries of $c \in \mathbf{R}^n$ were drawn i.i.d. from the uniform distribution over $[0, 1]$. Entries of $d \in \mathbf{Z}^n$ were drawn i.i.d. from the uniform distribution over the set $\{1, \ldots, 5\}$. We generated $b \in \mathbf{R}^m$ by first generating $\hat{z} \in \mathbf{Z}^n$, where each $\hat{z}_i$ was drawn from the uniform distribution over the set $\{0, \ldots, d_i\}$, and then setting $b = A\hat{z}$.

Results. We generated problem instances for a range of $10 \leq n \leq 100$. Figure 6 compares, for each n, the average value of $c^T z$ found by the NC-ADMM heuristic and by GUROBI over 10 instances. NC-ADMM was run from 10 random starting points for 100 iterations. The value of ρ for each starting point was drawn from the uniform distribution over $[0, 5]$. GUROBI's run time was limited to 10 minutes. NC-ADMM always found a feasible z with an objective value not much worse than that found by GUROBI. We also tried the relax-round-polish heuristic on the problem instances, but it never found a feasible z.

6.7 Maximum coverage problem

A collection of sets $\mathcal{S} = \{S_1, S_2, \ldots, S_m\}$ is defined over a domain of elements $\{e_1, e_2, \ldots, e_n\}$ with associated weights $\{w_i\}_{i=1}^n$. The goal is to find the collection of no more than k sets $\mathcal{S}' \subseteq \mathcal{S}$ that maximizes the total weight of elements covered by $\mathcal{S}'$ [KMN99, Hoc96, CK08]. Let $x_i \in \{0, 1\}$, for $i = 1, \ldots, n$, be a variable that takes 1 if element e_i is covered and 0 otherwise. Let $y \in \{0, 1\}^m$ be a variable with entry $y_j = 1$ if set j is selected. The problem

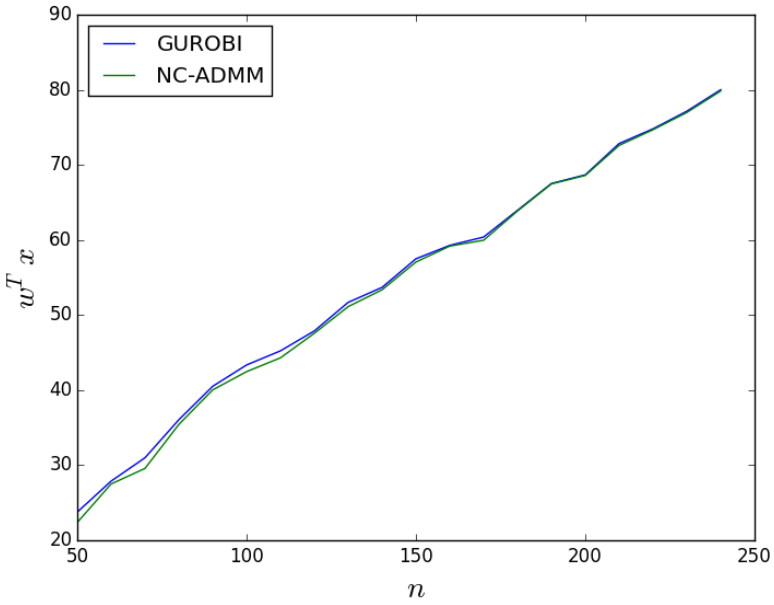

Figure 7: The average solution weight over 10 different instances.

is

$$\begin{array}{ll} \text{maximize} & w^T x \\ \text{subject to} & \sum_{j \in S_j} y_j \geq x_i, \quad i = 1, \ldots, n \\ & x_i \in \{0,1\}, \quad i = 1, \ldots, n \\ & y \in \{0,1\}^m \\ & \mathbf{card}(y) = k. \end{array} \tag{16}$$

Note that y is a Boolean vector with fixed cardinality.

Problem instances. We generated problems as follows. Each set contained each of the elements independently with a constant probability p. Hence the expected size of each set was np. There were $m = 3/p$ sets, so the expected total number of elements in all sets (with repetition) was equal to $mnp = 3n$. We set $k = 1/(3p)$. Each w_i was chosen uniformly at random from the interval $[0, 1]$.

Results. We generated problems as described above for $n = 50, 60, \ldots, 240$ and $p = 0.01$. For each value of n, we generated 10 problems and recorded the average weight $w^T x$ of the approximate solutions found by NC-ADMM and the optimal solutions found by GUROBI. Figure 7 shows the results of our comparison of NC-ADMM and GUROBI. Approximate solutions found by the relax-round-polish heuristic were far worse than those found by NC-ADMM for this problem.

6.8 Inexact graph isomorphism

Two (undirected) graphs are isomorphic if we can permute the vertices of one so it is the same as the other (i.e., the same pairs of vertices are connected by edges). If we describe them by their adjacency matrices A and B, isomorphism is equivalent to the existence of a permutation matrix $Z \in \mathbf{R}^{n\times n}$ such that $ZAZ^T = B$, or equivalently $ZA = BZ$.

Since in practical applications isomorphic graphs might be contaminated by noise, the inexact graph isomorphism problem is usually stated [ABK14, Ume88, CWH97], in which we want to find a permutation matrix Z such that the disagreement $\|ZAZ^T - B\|_F^2$ between the transformed matrix and the target matrix is minimized. Since $\|ZAZ^T - B\|_F^2 = \|ZA - BZ\|_F^2$ for any permutation matrix Z, the inexact graph isomorphism problem can be formulated as

$$\begin{array}{ll} \text{minimize} & \|ZA - BZ\|_F^2 \\ \text{subject to} & Z \in \mathcal{P}_n. \end{array} \tag{17}$$

If the optimal value of this problem is zero, it means that A and B are isomorphic. Otherwise, the solution of this problem minimizes the disagreement of ZAZ^T and B in the Frobenius norm sense.

Solving inexact graph isomorphism problems is of interest in pattern recognition [CFSV04, RP94], computer vision [SRS01], shape analysis [SKK04, HHW06], image and video indexing [Lee06], and neuroscience [VCP+11]. In many of the aforementioned fields graphs are used to represent geometric structures, and $\|ZAZ^T - B\|_F^2$ can be interpreted as the strength of geometric deformation.

Problem instances. It can be shown that if A and B are isomorphic and A has distinct eigenvalues and for all eigenvectors v of A for which $\mathbf{1}^T v \neq 0$, then the relaxed problem has a unique solution which is the permutation matrix that relates A and B [ABK14]. Hence, in order to generate harder problems, we generated the matrix A such that it violated these conditions. In particular, we constructed A for the Peterson graph (3-regular with 10 vertices), icosahedral graph (5-regular with 12 vertices), Ramsey graph (8-regular with 17 vertices), dodecahedral graph (3-regular with 20 vertices), and the Tutte-Coxeter graph (3-regular with 30 vertices). For each example we randomly permuted the vertices to obtain two isomorphic graphs.

Results. We ran NC-ADMM with 20 iterations and 5 restarts. For all of our examples NC-ADMM was able to find the permutation relating the two graphs. It is interesting to notice that running the algorithm multiple times can find different solutions if there is more than one permutation relating the two graphs. The relax-round-polish heuristic failed to find a solution for all of the aforementioned problems.

7 Conclusions

We have discussed the relax-round-polish and NC-ADMM heuristics and demonstrated their performance on many different problems with convex objectives and decision variables from a nonconvex set. Our heuristics are easy to extend to additional problems because they rely on a simple mathematical interface for nonconvex sets. We need only know a method for (approximate) projection onto the set. We do not require but benefit from knowing a convex relaxation of the set, a convex restriction at any point in the set, and the neighbors of any point in the set under some discrete distance metric. Adapting our heuristics to any particular problem is straightforward, and we have fully automated the process in the NCVX package.

We do not claim that our heuristics give state-of-the-art results for any particular problem. Rather, the purpose of our heuristics is to give a fast and reasonable solution with minimal tuning for a wide variety of problems. Our heuristics also take advantage of the tremendous progress in technology for solving general convex optimization problems, which makes it practical to treat solving a convex problem as a black box.

References

[ABK14] Y. Aflalo, A. Bronstein, and R. Kimmel. Graph matching: relax or not? *arXiv preprint arXiv:1401.7623*, 2014.

[ABK15] Y. Aflalo, A. Bronstein, and R. Kimmel. On convex relaxation of graph isomorphism. *Proceedings of the National Academy of Sciences*, 112(10):2942–2947, 2015.

[Ach09] T. Achterberg. SCIP: solving constraint integer programs. *Mathematical Programming Computation*, 1(1):1–41, 2009.

[AM85] W. N. Anderson and T. D. Morley. Eigenvalues of the Laplacian of a graph. *Linear and Multilinear Algebra*, 18(2):141–145, 1985.

[ApS15] MOSEK ApS. *The MOSEK optimization toolbox for MATLAB manual. Version 7.1 (Revision 28)*, 2015.

[BBC+08] P. Bonami, L. T. Biegler, A. R. Conn, G. Cornuéjols, I. E. Grossmann, C. D. Laird, J. Lee, A. Lodi, F. Margot, and N. Sawaya. An algorithmic framework for convex mixed integer nonlinear programs. *Discrete Optimization*, 5(2):186–204, 2008.

[Bec14] A. Beck. *Introduction to Nonlinear Optimization: Theory, Algorithms, and Applications with MATLAB*, volume 19. SIAM, 2014.

[Bel09] P. Belotti. Couenne: a user's manual. *Technical Report*, 2009.

[Bem15] A. Bemporad. Solving mixed-integer quadratic programs via nonnegative least squares. *5th IFAC Conference on Nonlinear Model Predictive Control*, pages 73–79, 2015.

[Ber14] D. P. Bertsekas. *Constrained optimization and Lagrange multiplier methods*. Academic press, 2014.

[BHA15] S. Boyd, M. Hast, and K. J. Astrom. MIMO PID tuning via iterated LMI restriction. *International Journal of Robust and Nonlinear Control*, 2015.

[BJS94] P. Brucker, B. Jurisch, and B. Sievers. A branch and bound algorithm for the job-shop scheduling problem. *Discrete applied mathematics*, 49(1):107–127, 1994.

[BM99] A. Bemporad and M. Morari. Control of systems integrating logic, dynamics, and constraints. *Automatica*, 35(3):407–427, 1999.

[BNW06] R. H. Byrd, J. Nocedal, and R. A. Waltz. Knitro: An integrated package for nonlinear optimization. In *Large-scale nonlinear optimization*, pages 35–59. Springer, 2006.

[Bol13] D. Boley. Local linear convergence of the alternating direction method of multipliers on quadratic or linear programs. *SIAM Journal on Optimization*, 23(4):2183–2207, 2013.

[BP12] A. Bemporad and P. Patrinos. Simple and certifiable quadratic programming algorithms for embedded linear model predictive control. In *Nonlinear Model Predictive Control*, volume 4, pages 14–20, 2012.

[BPC+11] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011.

[Bra10] A. M. Bradley. *Algorithms for the Equilibration of Matrices and their Application to Limited-Memory Quasi-Newton Methods*. PhD thesis, Stanford University, 2010.

[BRL01] A. Bemporad, J. Roll, and L. Ljung. Identification of hybrid systems via mixed-integer programming. In *IEEE Conference on Decision and Control*, volume 1, pages 786–792, 2001.

[BV04] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[CA96] J. M. Crawford and L. D. Auton. Experimental results on the crossover point in random 3-SAT. *Artificial intelligence*, 81(1):31–57, 1996.

[CA06] M. Carrión and J. M. Arroyo. A computationally efficient mixed-integer linear formulation for the thermal unit commitment problem. *IEEE Transactions on Power Systems*, 21(3):1371–1378, 2006.

[CB98] P. C. Chu and J. E. Beasley. A genetic algorithm for the multidimensional knapsack problem. *Journal of heuristics*, 4(1):63–86, 1998.

[CFSV04] D. Conte, P. Foggia, C. Sansone, and M. Vento. Thirty years of graph matching in pattern recognition. *International journal of pattern recognition and artificial intelligence*, 18(03):265–298, 2004.

[Cha12] R. Chartrand. Nonconvex splitting for regularized low-rank + sparse decomposition. *IEEE Transactions on Signal Processing*, 60(11):5810–5819, 2012.

[CK05] C. Chekuri and S. Khanna. A polynomial time approximation scheme for the multiple knapsack problem. *SIAM Journal on Computing*, 35(3):713–728, 2005.

[CK08] R. Cohen and L. Katzir. The generalized maximum coverage problem. *Information Processing Letters*, 108(1):15–22, 2008.

[CKP08] I. Castillo, F. J. Kampas, and J. D. Pintér. Solving circle packing problems by global optimization: numerical results and industrial applications. *European Journal of Operational Research*, 191(3):786–802, 2008.

[CPDB13] E. Chu, N. Parikh, A. Domahidi, and S. Boyd. Code generation for embedded second-order cone programming. In *Proceedings of the 2013 European Control Conference*, pages 1547–1552, 2013.

[CPL09] IBM ILOG CPLEX. User's manual for CPLEX, version 12.1. *International Business Machines Corporation*, 46(53):157, 2009.

[CPM10] J. P. S. Catalão, H. M. I. Pousinho, and V. M. F. Mendes. Scheduling of head-dependent cascaded hydro systems: Mixed-integer quadratic programming approach. *Energy Conversion and Management*, 51(3):524–530, 2010.

[CS03] C. R. Collins and K. Stephenson. A circle packing algorithm. *Computational Geometry*, 25(3):233–256, 2003.

[CW13] R. Chartrand and B. Wohlberg. A nonconvex ADMM algorithm for group sparsity with sparse groups. In *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6009–6013. IEEE, 2013.

[CWH97] A. D. J. Cross, R. C. Wilson, and E. R. Hancock. Inexact graph matching using genetic search. *Pattern Recognition*, 30(6):953–970, 1997.

[DBEY13] N. Derbinsky, J. Bento, V. Elser, and J. S. Yedidia. An improved three-weight message-passing algorithm. *arXiv:1305.1961 [cs.AI]*, 2013.

[DCB13] A. Domahidi, E. Chu, and S. Boyd. ECOS: An SOCP solver for embedded systems. In *Proceedings of the 12th European Control Conference*, pages 3071–3076. IEEE, 2013.

[DCB14] S. Diamond, E. Chu, and S. Boyd. CVXPY: A Python-embedded modeling language for convex optimization, version 0.2. `http://cvxpy.org/`, May 2014.

[DFJ54] G. Dantzig, R. Fulkerson, and S. Johnson. Solution of a large-scale traveling-salesman problem. *Journal of the operations research society of America*, 2(4):393–410, 1954.

[DM93] J. David and B. De Moor. The opposite of analytic centering for solving minimum rank problems in control and identification. In *Proceedings of the 32nd IEEE Conference on Decision and Control*, pages 2901–2902. IEEE, 1993.

[EB92] J. Eckstein and D. P. Bertsekas. On the Douglas-Rachford splitting method and the proximal point algorithm for maximal monotone operators. *Mathematical Programming*, 55(1-3):293–318, 1992.

[Ers14] T. Erseghe. Distributed optimal power flow using ADMM. *IEEE Transactions on Power Systems*, 29(5):2370–2380, 2014.

[FDM15] D. Frick, A. Domahidi, and M. Morari. Embedded optimization for mixed logical dynamical systems. *Computers and Chemical Engineering*, 72:21–33, 2015.

[FHT01] J. Friedman, T. Hastie, and R. Tibshirani. *The elements of statistical learning*, volume 1. Springer series in statistics Springer, Berlin, 2001.

[Fle73] R. Fletcher. An exact penalty function for nonlinear programming with inequalities. *Mathematical Programming*, 5(1):129–150, 1973.

[Flo95] C. A. Floudas. *Nonlinear and Mixed-Integer Optimization: Fundamentals and Applications*. Oxford University Press, 1995.

[GB14a] P. Giselsson and S. Boyd. Diagonal scaling in Douglas-Rachford splitting and ADMM. In *53rd Annual IEEE Conference on Decision and Control (CDC)*, pages 5033–5039, 2014.

[GB14b] P. Giselsson and S. Boyd. Monotonicity and restart in fast gradient methods. In *53rd Annual IEEE Conference on Decision and Control (CDC)*, pages 5058–5063, 2014.

[GB14c] P. Giselsson and S. Boyd. Preconditioning in fast dual gradient methods. In *53rd Annual IEEE Conference on Decision and Control (CDC)*, pages 5040–5045, 2014.

[GM12] S. Gualandi and F. Malucelli. Exact solution of graph coloring problems via constraint programming and column generation. *INFORMS Journal on Computing*, 24(1):81–100, 2012.

[GO15] Inc. Gurobi Optimization. Gurobi Optimizer Reference Manual, 2015.

[Gol70] M. Goldberg. The packing of equal circles in a square. *Mathematics Magazine*, pages 24–30, 1970.

[Gro07] J. L. Gross. *Combinatorial methods with computer applications*. CRC Press, 2007.

[GTSJ15] E. Ghadimi, A. Teixeira, I. Shames, and M. Johansson. Optimal parameter selection for the alternating direction method of multipliers (ADMM): Quadratic problems. *IEEE Transactions on Automatic Control*, 60(3):644–658, 2015.

[Hes69] M. R. Hestenes. Multiplier and gradient methods. *Journal of optimization theory and applications*, 4(5):303–320, 1969.

[HHW06] L. He, C. Y. Han, and W. G. Wee. Object recognition and recovery by skeleton graph matching. In *IEEE International Conference on Multimedia and Expo*, pages 993–996. IEEE, 2006.

[HL12] M. Hong and Z. Luo. On the linear convergence of the alternating direction method of multipliers. *arXiv preprint arXiv:1208.3922*, 2012.

[HLR14] M. Hong, Z. Luo, and M. Razaviyayn. Convergence analysis of alternating direction method of multipliers for a family of nonconvex problems. *arXiv preprint arXiv:1410.1390*, 2014.

[HM79] S. P. Han and O. L. Mangasarian. Exact penalty functions in nonlinear programming. *Mathematical programming*, 17(1):251–269, 1979.

[Hma10] H. Hmam. Quadratic optimization with one quadratic equality constraint. Technical report, Electronic Warfare and Radar Division, Defence Science and Technology Organisation (DSTO), Australia, 2010.

[Hoc96] D. S. Hochbaum. Approximating covering and packing problems: set cover, vertex cover, independent set, and related problems. In *Approximation algorithms for NP-hard problems*, pages 94–143. PWS Publishing Co., 1996.

[Hon14] M. Hong. A distributed, asynchronous and incremental algorithm for nonconvex optimization: An ADMM based approach. *arXiv preprint arXiv:1412.6058*, 2014.

[HPR13] K. L. Hoffman, M. Padberg, and G. Rinaldi. Traveling salesman problem. In *Encyclopedia of Operations Research and Management Science*, pages 1573–1578. Springer, 2013.

[HS16] K. Huang and N. D. Sidiropoulos. Consensus-ADMM for general quadratically constrained quadratic programming. *arXiv preprint arXiv:1601.02335*, 2016.

[JGR+14] J. L. Jerez, P. J. Goulart, S. Richter, G. Constantinides, E. C. Kerrigan, M. Morari, et al. Embedded online optimization for model predictive control at megahertz rates. *IEEE Transactions on Automatic Control*, 59(12):3238–3251, 2014.

[JMZ14] B. Jiang, S. Ma, and S. Zhang. Alternating direction method of multipliers for real and complex polynomial optimization models. *Optimization*, 63(6):883–898, 2014.

[Joh14] S. G. Johnson. The NLopt nonlinear-optimization package, 2014.

[Kal85] R. E. Kalman. Identification of noisy systems. *Russian Mathematical Surveys*, 40(4):25–42, 1985.

[KMN99] S. Khuller, A. Moss, and J. S. Naor. The budgeted maximum coverage problem. *Information Processing Letters*, 70(1):39–45, 1999.

[Kru56] J. B. Kruskal. On the shortest spanning subtree of a graph and the traveling salesman problem. *Proceedings of the American Mathematical society*, 7(1):48–50, 1956.

[Kuh05] H. W. Kuhn. The hungarian method for the assignment problem. *Naval Research Logistics (NRL)*, 52(1):7–21, 2005.

[Law85] E. L. Lawler. The traveling salesman problem: a guided tour of combinatorial optimization. *Wiley Series in Discrete Mathematics*, 1985.

[LB14] T. Lipp and S. Boyd. Variations and extensions of the convex-concave procedure, 2014.

[Lee06] J. Lee. A graph-based approach for modeling and indexing video data. In *IEEE International Symposium on Multimedia*, pages 348–355. IEEE, 2006.

[LP14] G. Li and T. K. Pong. Splitting methods for nonconvex composite optimization. *arXiv preprint arXiv:1407.0753*, 2014.

[LP15] G. Li and T. K. Pong. Global convergence of splitting methods for nonconvex composite optimization. *arXiv e-Print 1407.0753*, 2015.

[LS14] A. P. Liavas and N. D. Sidiropoulos. Parallel algorithms for constrained tensor factorization via the alternating direction method of multipliers. *arXiv preprint arXiv:1409.2383*, 2014.

[LW66] E. L. Lawler and D. E. Wood. Branch-and-bound methods: A survey. *Operations research*, 14(4):699–719, 1966.

[MB10] J. Mattingley and S. Boyd. Automatic code generation for real-time convex optimization. *Convex Optimization in Signal Processing and Communications*, pages 1–41, 2010.

[MB12] J. Mattingley and S. Boyd. CVXGEN: a code generator for embedded convex optimization. *Optimization and Engineering*, 13(1):1–27, 2012.

[Mer94] R. Merris. Laplacian matrices of graphs: a survey. *Linear algebra and its applications*, 197:143–176, 1994.

[MSL92] D. Mitchell, B. Selman, and H. Levesque. Hard and easy distributions of SAT problems. In *AAAI*, volume 92, pages 459–465, 1992.

[MWB11] J. Mattingley, Y. Wang, and S. Boyd. Receding horizon control: Automatic generation of high-speed solvers. *IEEE Control Systems Magazine*, 31(3):52–65, 2011.

[Neu37] J. Von Neumann. Some matrix inequalities and metrization of metric space. *Tomsk University Review*, 1:286–296, 1937.

[NF77] P. M. Narendra and K. Fukunaga. A branch and bound algorithm for feature subset selection. *IEEE Transactions on Computers*, 100(9):917–922, 1977.

[NTGTB15] L. Ning, T. Tryphon T. Georgiou, A. Tannenbaum, and S. Boyd. Linear models based on noisy data and the frisch scheme. *SIAM Review*, 57(2):167–197, 2015.

[OSB13] B. O'Donoghue, G. Stathopoulos, and S. Boyd. A splitting method for optimal control. *IEEE Transactions on Control Systems Technology*, 21(6):2432–2442, 2013.

[PB13] N. Parikh and S. Boyd. Proximal algorithms. *Foundations and Trends in Optimization*, 1(3):123–231, 2013.

[PB14] N. Parikh and S. Boyd. Block splitting for distributed optimization. *Mathematical Programming Computation*, 6(1):77–102, 2014.

[PCZ15] Z. Peng, J. Chen, and W. Zhu. A proximal alternating direction method of multipliers for a minimization problem with nonconvex constraints. *Journal of Global Optimization*, pages 1–18, 2015.

[PF07] L. G. Papageorgiou and E. S. Fraga. A mixed integer quadratic programming formulation for the economic dispatch of generators with prohibited operating zones. *Electric power systems research*, 77(10):1292–1296, 2007.

[PG89] G. Di Pillo and L. Grippo. Exact penalty functions in constrained optimization. *SIAM Journal on control and optimization*, 27(6):1333–1360, 1989.

[PR91] M. Padberg and G. Rinaldi. A branch-and-cut algorithm for the resolution of large-scale symmetric traveling salesman problems. *SIAM review*, 33(1):60–100, 1991.

[RP94] J. Rocha and T. Pavlidis. A shape analysis model with applications to a character recognition system. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(4):393–404, 1994.

[SA90] H. D. Sherali and W. P. Adams. A hierarchy of relaxations between the continuous and convex hull representations for zero-one programming problems. *SIAM Journal on Discrete Mathematics*, 3(3):411–430, 1990.

[SCPW12] J. Saunderson, V. Chandrasekaran, P. A. Parrilo, and A. S. Willsky. Diagonal and low-rank matrix decompositions, correlation matrices, and ellipsoid fitting. *SIAM Journal on Matrix Analysis and Applications*, 33(4):1395–1416, 2012.

[SKK04] T. B. Sebastian, P. N. Klein, and B. B. Kimia. Recognition of shapes by editing their shock graphs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(5):550–571, 2004.

[Slu69] A. V. D. Sluis. Condition numbers and equilibration of matrices. *Numerische Mathematik*, 14(1):14–23, 1969.

[SLY+14] W. Shi, Q. Ling, K. Yuan, G. Wu, and W. Yin. On the linear convergence of the admm in decentralized consensus optimization. *IEEE Transactions on Signal Processing*, 62(7):1750–1761, 2014.

[SM99] R. A. Stubbs and S. Mehrotra. A branch-and-cut method for 0-1 mixed convex programming. *Mathematical programming*, 86(3):515–532, 1999.

[SP14] S.You and Q. Peng. A non-convex alternating direction method of multipliers heuristic for optimal power flow. In *2014 IEEE International Conference on Smart Grid Communications (SmartGridComm)*, pages 788–793. IEEE, 2014.

[Spe13] E. Specht. Packomania. http://www.packomania.com/, October 2013.

[Spi85] J. E. Spingarn. Applications of the method of partial inverses to convex programming: decomposition. *Mathematical Programming*, 32(2):199–223, 1985.

[SRS01] C. Schellewald, S. Roth, and C. Schnörr. Evaluation of convex optimization techniques for the weighted graph-matching problem in computer vision. In *Pattern Recognition*, pages 361–368. Springer, 2001.

[Ste05] K. Stephenson. *Introduction to circle packing: The theory of discrete analytic functions*. Cambridge University Press, 2005.

[TMBB15] R. Takapoui, N. Moehle, S. Boyd, and A. Bemporad. A simple effective heuristic for embedded mixed-integer quadratic programming. *arXiv preprint arXiv:1509.08416*, 2015.

[TS05a] M. Tawarmalani and N. V. Sahinidis. A polyhedral branch-and-cut approach to global optimization. *Mathematical Programming*, 103:225–249, 2005.

[TS05b] M. Tawarmalani and N. V. Sahinidis. A polyhedral branch-and-cut approach to global optimization. *Mathematical Programming*, 103(2):225–249, 2005.

[Ull11] F. Ullmann. FiOrdOs: A Matlab toolbox for C-code generation for first order methods. Master's thesis, ETH Zurich, 2011.

[Ume88] S. Umeyama. An eigendecomposition approach to weighted graph matching problems. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10(5):695–703, 1988.

[VCP+11] J. T. Vogelstein, J. M. Conroy, L. J. Podrazik, S. G. Kratzer, E. T. Harley, D. E. Fishkind, R. J. Vogelstein, and C. E. Priebe. Large (brain) graph matching via fast approximate quadratic programming. *arXiv preprint arXiv:1112.5507*, 2011.

[WB10] Y. Wang and S. Boyd. Fast model predictive control using online optimization. *IEEE Transactions on Control Systems Technology*, 18(2):267–278, 2010.

[WLMK09] A. Waechter, C. Lairdl, F. Margot, and Y. Kawajir. Introduction to IPOPT: A tutorial for downloading, installing, and using IPOPT, 2009.

[WN99] S. J. Wright and J. Nocedal. *Numerical Optimization*, volume 2. Springer New York, 1999.

[WXX14] F. Wang, Z. Xu, and H. Xu. Convergence of alternating direction method with multipliers for non-convex composite problems. *arXiv preprint arXiv:1410.8625*, 2014.

[XYWZ12] Y. Xu, W. Yin, Z. Wen, and Y. Zhang. An alternating direction algorithm for matrix completion with nonnegative factors. *Frontiers of Mathematics in China*, 7(2):365–384, 2012.